# vSphere Data Protection Administration Guide

vSphere Data Protection 6.1

This document supports the version of each product listed and supports all subsequent versions until the document is replaced by a new edition. To check for more recent editions of this document, see http://www.vmware.com/support/pubs.

EN-001795-00

vmware®

You can find the most up-to-date technical documentation on the VMware Web site at:

http://www.vmware.com/support/

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com



**VMware, Inc.**
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

# Contents


**1** Understanding vSphere Data Protection 13

Introduction to vSphere Data Protection 14
Benefits of VDP 14
VDP Functionality 15
Image Level Backup and Restore 16
vCenter Server Backup and Restore 16
Single VMDK Backup and Restore 16
Guest-level Backup and Restore 17
Replication 17
File Level Recovery 17
Customer Experience Improvement Program 17
VDP Architecture 17

**2** VDP Installation and Configuration 21

vSphere Data Protection Capacity Requirements 22
Software Requirements 22
Hardware Versions and Migration 22
vSphere Hosts and vSphere Flash Read Cache Compatibility and Performance 22
Unsupported Disk Types 22
Unsupported Virtual Volumes 23
System Requirements 23
VDP System Requirements 23
IPv6 Requirements 23
Preinstallation Configuration 23
DNS Configuration 23
NTP Configuration 24
vCenter Hosts and Clusters View 24
User Account Configuration 24
VDP Best Practices 26
General Best Practices 26
Deployment Best Practices 26
HotAdd Best Practices 27
Storage Capacity for Initial VDP Deployment 27
Sizing 27
Monitoring VDP Capacity 28
VDP Installation 29
Deploying the OVF Template 29
Initial Configuration 30

**3** VDP Appliance Migration 35

Migrating the VDP Appliance from 5.8.x or Later to 6.1 36

**4** Post-Installation Configuration of VDP Appliance 37

About the VDP Configure Utility 38
Viewing Status 38
Starting and Stopping Services 39
Collecting VDP Logs or Diagnostic Information 40

Modifying Configuration Settings 42
Network Settings 42
vCenter Server Registration 42
Rolling Back an Appliance 43
Emergency Restore 44
Automatic Host Detection 46
Refreshing Restore Points 46
Reconnecting the Host to the vCenter 46

**5** Securing Communication between VDP and vCenter 47
Replacing the VDP Certificate 48
Authenticating the VDP Server for Secured Communications 49
Securing the VDP Appliance Communication with the vCenter Server 50
Securing the Proxy Communication with the vCenter Server 51

**6** Configuring VDP 53
Login Security 54
Configuring the Customer Experience Improvement Program 54
Configuring and Monitoring 55
Viewing Backup Appliance Configuration 55
Editing the Backup Window 57
Configuring the Email Notifications and Reports 57
Viewing the User Interface Log 59
Running an Integrity Check 59
Monitoring VDP Activity 60
Viewing Recent Tasks 60
Viewing Alarms 61
Viewing the Event Console 62
Persisted and Pop-up Messages to Show the Capacity Usage Issue 62
VDP Shutdown and Startup Procedures 62

**7** Proxies 63
Proxy Overview 64
Considerations Before Deploying an External Proxy 64
Deployment of External Proxies 64
Number of Proxies to Deploy and Proxy Throughputs per Proxy 64
Managing Proxy Throughput 65
External Proxy Support 66
Adding an External Proxy 67
Disabling the Internal Proxy 68
(Optional) Configuring Proxy Certificate Authentication 68
Monitoring External Proxy Health Status 69
Health Status Criteria 69
External Proxy Logs 69

**8** Storage Management 71
Creating New Storage 72
Minimum Storage Performance 73
Attaching Existing VDP Disks 74
Detaching and Reattaching Storage 75
Viewing the Storage Configuration 76
Enabling the Seek Test 77
Changing the Configuration File to Enable the Seek Test 77

**9** Data Domain Integration 79

VDP and Data Domain System Integration 80
Architecture Overview 80
VDP Client Support 80
Best Practices 81
  Data Domain Limitations 81
  Backup 81
  Restore 82
  Security - Encryption 82
  Data Migration 82
Pre-integration Requirements 82
  Network Throughput 82
  Network Configuration 83
  NTP Configuration 83
  Licensing 83
  Port Usage and Firewall Requirements 83
  Capacity 83
  Data Domain System Streams 84
  Existing Backup Products in Use with Data Domain 84
Preparing the Data Domain system for VDP Integration 84
Adding a Data Domain System 85
Editing the Data Domain System 87
Deleting the Data Domain System from the VDP Appliance 87
Backups with VDP and Data Domain 89
  How Backups Work with VDP and Data Domain 89
  Where Backup Data is Stored 89
  How VDP Appliance Manages Backup Data 89
  Supported Backup Types 90
  Canceling and Deleting Backups 90
  Selecting a Data Domain Target for Backups 90
Replication Control 90
  Replication Data Flow 91
  Replication Schedule 91
  Replication Configuration 91
  Replication Monitoring with VDP 91
Server Maintenance Activity Monitoring 91
Restoring Avamar Checkpoint backups from Data Domain systems 92
  Assumptions for the Restore Operation 92
  Performing the Checkpoint Restore 92
Monitoring Data Domain from the VDP Appliance 94
  Monitoring by Using the vSphere Web Client 94
  Monitoring by Using the VDP Configuration Utility 95
Reclaiming Storage on a Full Data Domain System 95
Common Problems and Solutions 97
  Backup Fails if the Data Domain System is Offline 97
  Rolling Back after Deleting a Data Domain System 97

**10** VDP Disk Expansion 99

Prerequisites 100
VMFS Heap Size Recommendations 100
Performing Disk Expansion 101
  Viewing the Storage Configuration 102
Performance Analysis 103
  Running the Performance Analysis Test 103

Disk Expansion with Essentials Plus 103

**11** Using VDP 105

Accessing VDP 106

Accessing the VDP Appliance by Using CLI 106

Useful Commands 106

Understanding the VDP User Interface 107

Switching VDP Appliances 108

VDP User Interface 108

Viewing Information from the Reports Tab 108

Refresh 108

Task Failures Tab 109

Job Details Tab 110

Unprotected Clients Tab 110

**12** Managing Backups 113

Backup Jobs 114

Choosing the Virtual Machines 114

Identifying retired virtual machines 114

Specifying the Backup Schedule 115

Setting the Retention Policy 115

Creating a Full Image Backup Job 116

Creating a Backup Job on Individual Disks 117

Unsupported disk types 118

Limitation 118

Migration on Individual Disks 119

Viewing Status and Backup Job Details 119

Editing a Backup Job 119

Cloning a Backup Job 119

Deleting a Backup Job 119

Enabling or Disabling a Backup Job 120

Running Existing Backup Jobs Immediately 120

Locking and Unlocking a Backup 121

Migrating Backup Jobs from VDP to Avamar 121

Guidelines 121

Recommendations 122

Prerequisites 122

Procedure 122

Troubleshooting 123

**13** Automatic Backup Verification 125

About Automatic Backup Verification 126

Limitations 126

Best Practices 126

Creating a New Backup Verification Job 127

Editing a Backup Verification Job 129

Cloning a Backup Verification Job 129

Executing and Running a Backup Verification Job 129

Monitoring Backup Verification 130

Enabling and Disabling a Backup Verification Job 130

Deleting a Backup Verification Job 130

**14** Managing Restores 131
- Restore Operations 132
  - Limitations 132
- Selecting Backups to Restore 132
- Filtering for List of Backups 133
- Restores when Snapshots Are Present 133
- Restoring Image Backups to the Original Location 133
- Restoring Image Backups to a New Location 135
  - Disabling the vMotion Feature Before Performing Restores to a New Location 136
- Restoring Backups to Individual SCSI Disks 136
- Deleting a Backup from the Restore Tab 137
- Clearing all Selected Backups from the Restore Tab 138

**15** Replication 139
- Replication Jobs 140
  - Replication Compatibility 140
  - Replication and Data Domain 142
  - Best Practices for Replication 142
  - Limitations 143
  - Defining Backup Types for a Replication Job 143
- Creating a Replication Job 143
- Managing Destinations 147
- Editing a Replication Job 147
- Cloning a Replication Job 147
- Deleting a Replication Job 148
- Enabling or Disabling a Replication Job 148
- Viewing Status and Replication Job Details 148
- Running Existing Replication Jobs Immediately 148
- Replication Back to the Source 148
  - Node Structure for Recovered Backups 148
  - Node Structure of Backups Replicated Again 148
  - Replication Destinations 149
- Replication Recovery Compatibility 149
- Enabling or Disabling Replication Recovery 150
- Replication Recovery 150
- Multi-Tenancy 151

**16** Using File Level Restore 153
- Introduction to the VDP Restore Client 154
  - LVM and EXT Support 154
  - File Level Restore Limitations 154
  - Unsupported VMDK Configurations 155
  - Unsupported Windows Configurations 155
- Logging in to the Restore Client 155
  - Basic Login 155
  - Advanced Login 156
- Mounting Backups 157
- Filtering Backups 157
- Navigating Mounted Backups 157
- Performing File Level Restores 157
  - Using the Restore Client in Basic Login Mode 157
  - Using the Restore Client in Advanced Login Mode 158
- Monitoring Restores 159

**17** VDP Application Support 161
VDP Application Support 162
Installing Application Agents 162
Checking the User Account Control Setting on Microsoft Windows 162
Installing VDP Clients when UAC is Enabled 162
Backing Up and Restoring Microsoft SQL Servers 163
Microsoft SQL Server Options 163
Hardware Requirements 163
Authentication Requirements 164
Microsoft SQL Server Support 164
Installing VDP for SQL Server Client 165
Configuring the Cluster Client in a Failover Cluster 166
Configuring the Cluster Client for an AlwaysOn Availability Group 167
Creating Backup Jobs for Microsoft SQL Servers 168
Restoring Backups of Microsoft SQL Servers 171
Monitoring Client Activity 172
Uninstalling the VDP Plug-in for SQL Server 172
Backing Up and Restoring Microsoft Exchange Servers 172
Microsoft Exchange Server Options 173
Microsoft Exchange Server Support 173
Microsoft .NET Framework 4 Requirement 173
Hardware Requirements 173
Unsupported Microsoft Exchange Servers 173
Installing VDP for Exchange Server Client 174
Installing in a DAG or Cluster Environment 174
Configuring an Exchange DAG Client 175
Using the VMware Exchange Backup User Configuration Tool 176
Manually Configuring the VDP Backup Service 178
Creating Backup Jobs for Microsoft Exchange Servers 178
Restoring Backups of Microsoft Exchange Servers 180
Suspending Replication in a DAG or Cluster 181
Monitoring Client Activity 181
Uninstalling the Exchange Server Plug-in 181
Granular Level Recovery on Microsoft Exchange Servers 182
Backing Up and Restoring Microsoft SharePoint Servers 185
Hardware Requirements 185
Supported Microsoft SharePoint Servers 185
Installing VDP for SharePoint Server Client 186
Creating Backup Jobs for Microsoft SharePoint Servers 187
Restoring Backups of Microsoft SharePoint Servers 188
Monitoring Client Activity 188
Uninstalling the VDP Plug-in for SharePoint Server 188

**18** VDP Disaster Recovery 189
Basic Disaster Recovery 190

**A** VDP Port Usage 191

**B** Minimum Required vCenter User Account Permissions 195

**C** VDP Troubleshooting 199
Troubleshooting the VDP Appliance Installation Issues 200
Troubleshooting the Installer Package Issues 200
Troubleshooting the VDP Management Issues 200
Troubleshooting the VDP Backup Issues 201

Troubleshooting the VDP Backup Performance Issues 202
Troubleshooting the VDP Restore Issues 203
Troubleshooting the VDP Replication Issues 204
Troubleshooting the VDP Integrity Check Issues 205
Troubleshooting the Automatic Backup Verification Issues 205
Troubleshooting the Restore Client (File Level Recovery) Issues 206
Troubleshooting the VDP Licensing Issues 207
Troubleshooting the VDP Appliance Issues 208
Troubleshooting the VDP Microsoft Exchange Server Issues 208
Troubleshooting the VDP Microsoft SQL Server Issues 210
Troubleshooting the VDP Microsoft SharePoint Server Issues 210
Troubleshooting the Storage Capacity Issues 210
Monitoring the VDP Local Storage Capacity 211
Monitoring the Data Domain Storage Capacity 211
Monitoring the Capacity Issues 211
General Steps to Free Storage Space 212
Accessing the VDP Knowledge Base Articles 212

Index 213

# About This Book

The *vSphere Data Protection Administration Guide* describes how to install and manage backups for small and medium businesses. This guide also includes troubleshooting scenarios and recommendations for resolution.

## Intended Audience

This book is for anyone who wants to provide backup solutions by using vSphere® Data Protection (VDP). The information in this book is for experienced Windows or Linux system administrators who are familiar with virtual machine technology and datacenter operations.

## Typographical conventions

VMware uses the following type style conventions in this document:

| | |
|---|---|
| **Bold** | Use for names of interface elements, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks) |
| *Italic* | Use for full titles of publications referenced in text |
| `Monospace` | Use for:<br>■ System output, such as an error message or script<br>■ System code<br>■ Pathnames, filenames, prompts, and syntax<br>■ Commands and options |
| *`Monospace italic`* | Use for variables |
| **`Monospace bold`** | Use for user input |
| [ ] | Square brackets enclose optional values |
| \| | Vertical bar indicates alternate selections , that is, **or** |
| { } | Braces enclose content that the user must specify, such as x or y or z |
| ... | Ellipses indicate nonessential information omitted from the example |

## VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to http://www.vmware.com/support/pubs.

# Document Feedback

VMware welcomes your suggestions for improving our documentation. Send your feedback to docfeedback@vmware.com.

# Technical Support and Education Resources

The following sections describe the technical support resources available to you. To access the current versions of other VMware books, go to http://www.vmware.com/support/pubs.

## Online Support

To use online support to submit technical support requests, view your product and contract information, and register your products, go to http://www.vmware/support/phone_support.html.

## Support Offerings

To find out how VMware support offerings can help meet your business needs, go to http://www.vmware.com/support/services.

## VMware Professional Services

VMware Education Services courses offer extensive hands-on labs, case study examples, and course materials designed to be used as on-the-job reference tools. Courses are available onsite, in the classroom, and live online. For onsite pilot programs and implementation best practices, VMware Consulting Services provides offerings to help you assess, plan, build, and manage your virtual environment. To access information about education classes, certification programs, and consulting services, go to http://www.vmware.com/services.

# Understanding vSphere Data Protection

1

This chapter includes the following topics:


- "Introduction to vSphere Data Protection" on page 14
- "Benefits of VDP" on page 14
- "VDP Functionality" on page 15
- "Replication" on page 17
- "File Level Recovery" on page 17
- "Customer Experience Improvement Program" on page 17
- "VDP Architecture" on page 17

# Introduction to vSphere Data Protection

vSphere Data Protection (VDP) is a robust, simple to deploy, disk-based backup and recovery solution that is powered by EMC. VDP is fully integrated with the VMware vCenter server and enables centralized and efficient management of backup jobs while storing backups in deduplicated destination storage locations.

The VMware vSphere web client interface is used to select, schedule, configure, and manage backups and recoveries of virtual machines.

During a backup, VDP creates a quiesced snapshot of the virtual machine. Deduplication is automatically performed with every backup operation.

The following terms are used throughout this document in the context of backup and recovery.

- A **datastore** is a virtual representation of a combination of underlying physical storage resources in the datacenter. A datastore is the storage location (for example, a physical disk, a RAID, or a SAN) for virtual machine files.
- **Changed Block Tracking (CBT)** is a VMkernel feature that keeps track of the storage blocks of virtual machines as they change over time. The VMkernel keeps track of block changes on virtual machines, which enhances the backup process for applications that have been developed to take advantage of VMware's vStorage APIs.
- **File Level Recovery (FLR)** allows local administrators of protected virtual machines to browse and mount backups for the local machine. From these mounted backups, the administrator can then restore individual files. FLR is accomplished using the VDP Restore Client.
- **VMware vStorage APIs for Data Protection (VADP)** enables backup software to perform centralized virtual machine backups without the disruption and overhead of running backup tasks from inside each virtual machine.
- **Virtual Machine Disk (VMDK)** is a file or set of files that appears as a physical disk drive to a guest operating system. These files can be on the host machine or on a remote file system.
- **VDP appliance** is a purpose-built virtual appliance for VDP.

# Benefits of VDP

The benefits of vSphere Data Protection (VDP) are as follows:

- Provides fast and efficient data protection for all of your virtual machines, even those powered off or migrated between vSphere hosts.
- Significantly reduces disk space consumed by backup data using patented variable-length deduplication across all backups.
- Reduces the cost of backing up virtual machines and minimizes the backup window by using Change Block Tracking (CBT) and VMware virtual machine snapshots.
- Allows for easy backups without the need for third-party agents installed in each virtual machine.
- Uses a simple, straight-forward installation as an integrated component within vSphere, which is managed by a web portal.
- Provides direct access to VDP configuration integrated into the vSphere web client.
- Protects backups with checkpoint and rollback mechanisms.
- Provides simplified recovery of Windows and Linux files with end-user initiated file-level recoveries from a web-based interface.
- By using emergency restore, VDP provides a method to restore virtual machines when either the vCenter server is unavailable or the user is unable to access the VDP user interface by using the vSphere web client.
- Through replication, enables you to avoid data loss if the source VDP appliance fails because copies of the backups are available on a destination target.

- Deduplication Store Benefits

  Enterprise data is highly redundant, with identical files or data stored within and across systems (for example, OS files or documents sent to multiple recipients). Edited files also have tremendous redundancy with previous versions. Traditional backup methods magnify this by storing all of the redundant data over and over again. vSphere Data Protection uses patented deduplication technology to eliminate redundancy at both the file and the subfile data segment level.

- Variable vs. Fixed-Length Data Segments

  A key factor in eliminating redundant data at a segment (or subfile) level is the method for determining segment size. Fixed-block or fixed-length segments are commonly employed by snapshot and some deduplication technologies. Unfortunately, even small changes to a dataset (for example, inserting data at the beginning of a file) can change all fixed-length segments in a dataset, despite the fact that very little of the dataset has been changed. vSphere Data Protection uses an intelligent variable-length method for determining segment size that examines the data to determine logical boundary points, which increases efficiency.

- Logical Segment Determination

  VDP uses a patented method for segment size determination designed to yield optimal efficiency across all systems. VDP's algorithm analyzes the binary structure of a dataset to determine segment boundaries that are context-dependent. Variable-length segments average 24 KB in size and are further compressed to an average of 12 KB. By analyzing the binary structure within the VMDK file, VDP works for all file types and sizes and deduplicates the data.

# VDP Functionality

The vSphere Data Protection product, starting with version 6.0, includes all features that were previously included in VDP Advanced. VDP functionality is included as part of vSphere Essential+ and does not require a specific license key. The following table lists VDP functionality.

**Table 1-1.** VDP Functionality

| Feature | VDP |
|---|---|
| Virtual machines supported per VDP appliance | Up to 400 |
| Number of appliances supported per vCenter | Up to 20 |
| Available storage size | 0.5 TB, 1 TB, 2 TB, 4 TB, and 8 TB |
| Support for image level backups | Yes |
| Support for individual disk backups | Yes |
| Support for image level restore jobs | Yes |
| Support for image level replication jobs | Yes |
| Support for direct to host recovery | Yes |
| Support for detachable/remountable data partitions | Yes |
| Support for file level recovery (FLR) | Yes, supports LVM and EXT4 with external proxies |
| Support for guest-level backups and restores of Microsoft Exchange Servers, SQL Servers, and SharePoint Servers | Yes |
| Support for application-level replication | Yes |
| Support for backing up to a Data Domain system | Yes |
| Ability to restore to a granular level on Microsoft Servers | Yes |
| Support for automatic backup verification (ABV) | Yes |
| Support for external proxies | Yes, up to 24 simultaneous virtual machines if the maximum number of 8 external proxies are deployed. |
| Support for Customer Experience Improvement Program | Yes |

## Image Level Backup and Restore

vSphere Data Protection creates image level backups, which are integrated with the vStorage API for Data Protection, a feature set within vSphere to offload the backup processing overhead from the virtual machine to the VDP appliance. The VDP appliance communicates with the vCenter server to make a snapshot of a virtual machine's .vmdk files. Deduplication takes place within the appliance by using a patented variable-length deduplication technology.

To support the large scale and continually expanding size of many VMware environments, each VDP appliance can simultaneously back up to 8 virtual machines if the internal proxy is used, or back up to 24 virtual machines if the maximum number of 8 external proxies are deployed with the VDP appliance.

To increase the efficiency of image level backups, VDP utilizes the Changed Block Tracking (CBT) feature, which greatly reduces the backup time of a given virtual machine image and provides the ability to process a large number of virtual machines within a particular backup window.

By leveraging CBT during restores, VDP offers fast and efficient recoveries of virtual machines to their original location. During a restore process, VDP uses CBT to determine which blocks have changed since the last backup. The use of CBT reduces data transfer within the vSphere environment during a recovery operation and more importantly reduces the recovery time.

Additionally, VDP automatically evaluates the workload between both restore methods (full image restore or a recovery leveraging CBT) and performs the method resulting in the fastest restore time. This is useful in scenarios where the change rate since the last backup in a virtual machine being restored is very high and the overhead of a CBT analysis operation would be more costly than a direct full-image recovery. VDP determines which method results in the fastest image recovery times for virtual machines in the environment.

## vCenter Server Backup and Restore

VDP supports backups of a vCenter Server by using an embedded Platform Service Controller and an external Platform Service Controller. VDP can perform file system quiescing during backups.

You must create a separate backup job that contains only vCenter Server. Schedule this backup job in the VDP backup window when the vCenter Server's nd VDP's workloads are the least to ensure that you can generate a file system-consistent snapshot during the backup. If the vCenter Server's workload is heavy, the backup can fail because of failure to quiesce the VM. vCenter Server is also reasonably resilient to crash-consistent restore. However, VMware supports backup and restore of vCenter Server by using VDP on the best-effort basis, and does not guarantee a successful restore.

Section 2.5.e in https://www.vmware.com/files/pdf/support/support_terms_conditions.pdf provides information about the terms and conditions for the customers' responsibilities to protect data.

To perform backups of vCenter Server by using an external Platform Services Controller, perform the steps that http://kb.vmware.com/kb/2110294 describes.

## Single VMDK Backup and Restore

A full image backup job includes all disks in the entire virtual machine (VM) in a single image backup. Individual disk backup jobs allow you to select only the disks you need. An image level backup of a VM with unsupported disk types does not include the unsupported disk types because of snapshot limitations.

When you restore a VM, the VDP appliance restores the VM configuration file (.vmx), which results in the creation of all VMDKs from the original VM. If any of the original VMDKs were not backed up, the restore process creates them as provisional VMDKs. The VM may not be fully functional in this case. The protected VMDKs, however, can be accessed from the restore.

See "Creating a Backup Job on Individual Disks" on page 117 for instructions on backing up individual disks.

### Guest-level Backup and Restore

VDP supports guest-level backups for Microsoft SQL Servers, Exchange Servers, and Share Point Servers. With guest-level backups, client agents (VMware VDP for SQL Server Client, VMware VDP for Exchange Server Client, or VMware VDP for SharePoint Server Client) are installed on the SQL Server, Exchange Server, or SharePoint Server in the same manner that backup agents are typically installed on physical servers.

The advantages of VMware guest-level backups are:

- Provides additional application support for Microsoft SQL Server, Microsoft Exchange Server, or SharePoint Server inside the virtual machines
- Support for backing up and restoring entire Microsoft SQL Server, Microsoft Exchange Server, or SharePoint Servers or selected databases
- Identical backup methods for physical and virtual machines

See “VDP Application Support” on page 161 for additional information on guest-level backup and restore.

## Replication

Replication enables you to avoid data loss if the source VDP appliance fails because copies of the backups are available on the destination target.

Replication jobs determine which backups are replicated, and when and to where the backups are replicated. With scheduled or ad hoc replication jobs for clients that have no restore points, only the client is replicated on the destination server. Backups created with VDP 6.0 or later can be replicated to another VDP appliance, to an EMC Avamar server, or to a Data Domain system. If the target VDP appliance is 5.8 or earlier, then the target must be VDP Advanced or Replication Target Identity.

See Chapter 15, “Replication,” on page 139 for additional information on Replication.

## File Level Recovery

File Level Recovery (FLR) allows local administrators of protected virtual machines to browse and mount backups for the local machine. From these mounted backups, the administrator can then restore individual files. FLR is accomplished by using the VDP Restore Client.

See Chapter 16, “Using File Level Restore,” on page 153 for additional information on FLR.

## Customer Experience Improvement Program

The Customer Experience Improvement Program is an option that enables you to send encrypted configuration and usage information about the VDP environment to VMware servers for analysis. The purpose of the Customer Experience Improvement Program is to help VDP improve the quality, reliability, and functionality of the VDP product. The Customer Experience Improvement Program is not enabled by default. During installation, you can enable the Customer Experience Improvement Program from the Product Improvement page in the VDP Configure Utility. You can also enable or disable this option any time after the installation of VDP from the post-installation UI. Refer to “Configuring the Customer Experience Improvement Program” on page 54 for more information.

## VDP Architecture

VDP can be deployed to any storage supported by vSphere. Supported storage includes VMFS, NFS, and VSAN datastores. Management of VDP is performed by using the vSphere web client.

Backup data is deduplicated and stored in the .vmdk files that make up the VDP virtual appliance or a supported Data Domain appliance.

The following figure illustrates the general architecture of VDP:

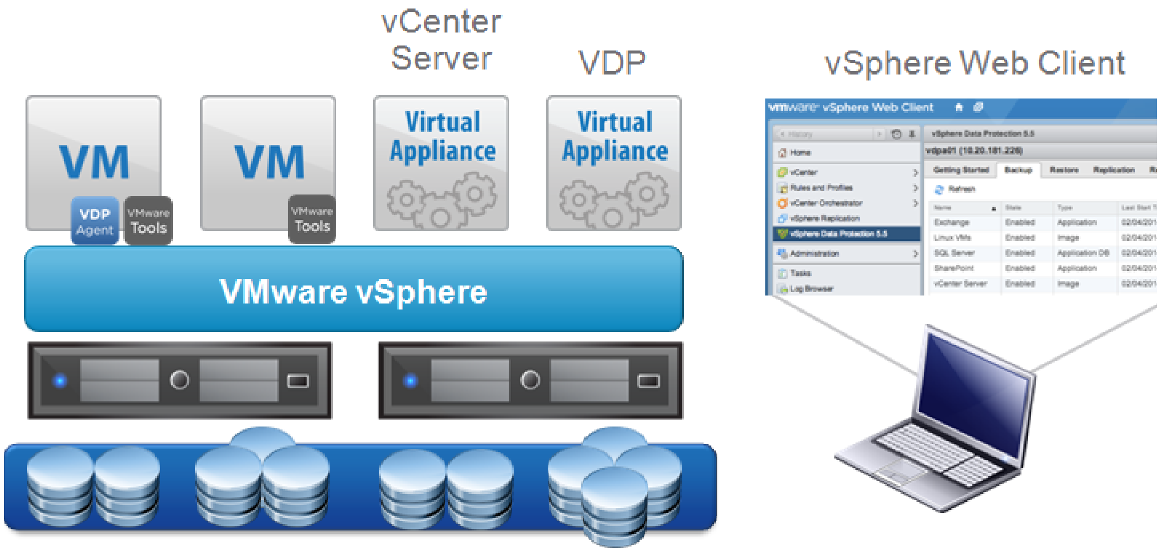



**Figure 1-1.** VDP general architecture

VDP general architecture consists of the following components:

- vCenter server 5.5 or later

  vCenter server 6.0 or later for VDP 6.1.6 or later

- VDP virtual appliance that is installed on any vSphere host version 5.1 or later
- vSphere web client
- Application backup agents

The following figures illustrate the detailed architecture of VDP and the legend of the architecture:

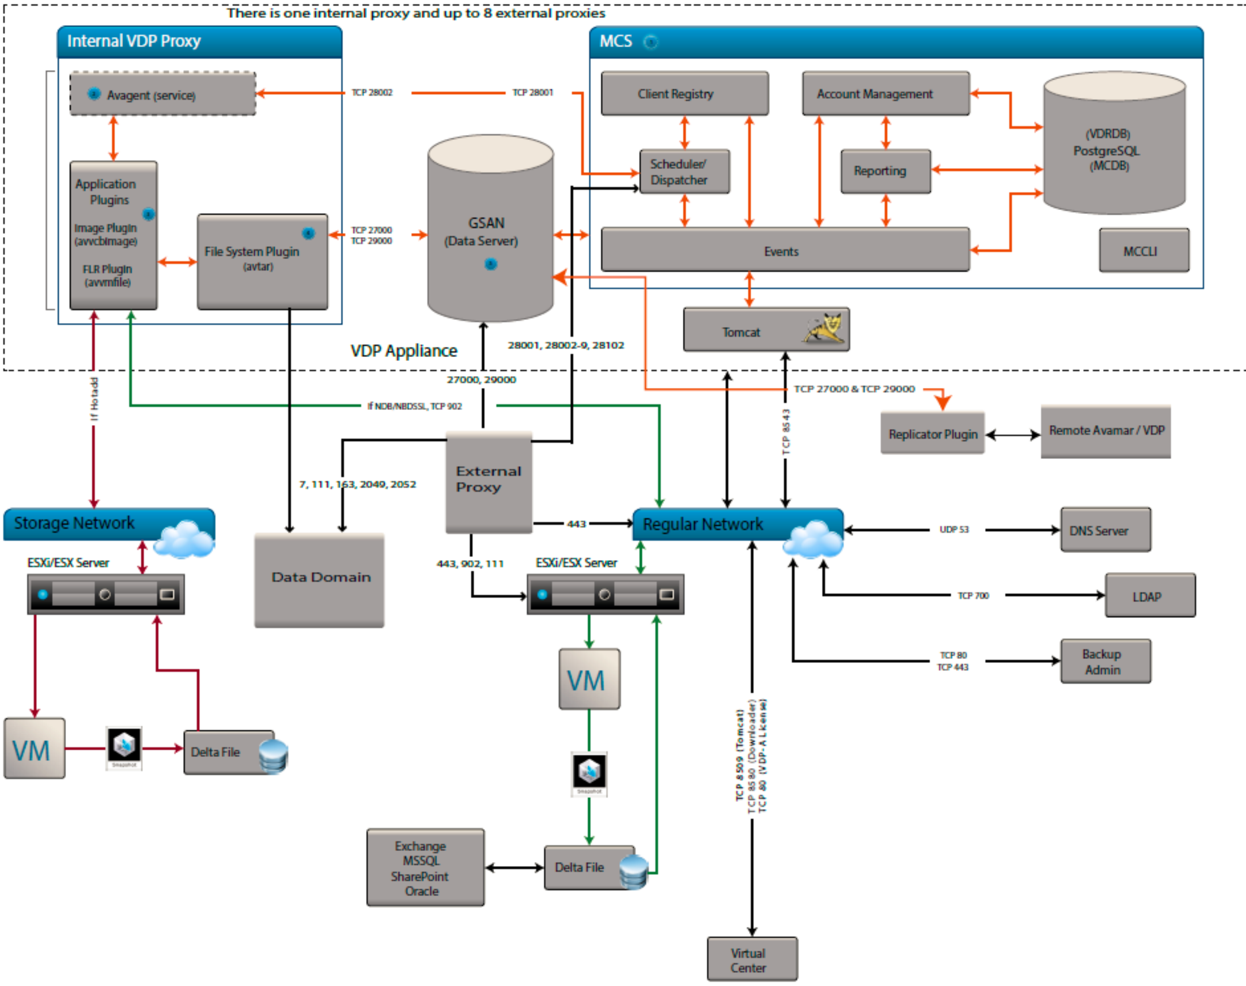



**Figure 1-2.** VDP detailed architecture

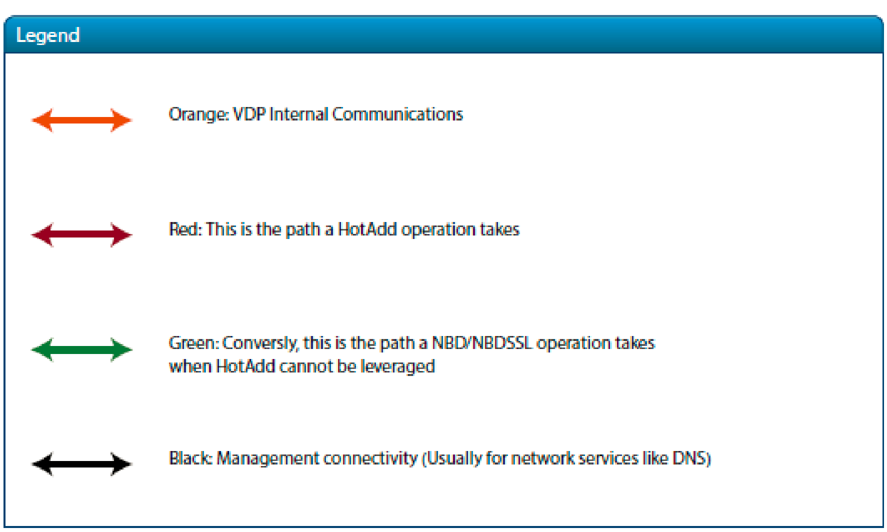



**Figure 1-3.** Legend of the VDP detailed architecture

The VDP detailed architecture consists of the following important components:

- **MCS**: Management Console Server (MCS) provides centralized administration such as scheduling, monitoring, and management for the VDP server, and runs the server-side processes that the administrator uses. The processes include all the Java processes and postgres (SQL) processes.

  The corresponding log file for this component is **/usr/local/avamar/var/mc/server_log/mcserver.log**.

- **AvAgent**: Runs as a service within each proxy on the VDP appliance. It establishes and maintains communication with MCS. AvAgent polls the MCS server every 15 seconds for incoming work orders. In response to a work order, such as a backup or a restore, AvAgent spawns the AvVcbImage which in turn calls AvTar.

  The corresponding log file for this component is **/usr/local/avamarclient/var-proxy-N/avagent.log**.

- **AvVcbImage**: Enables browsing, backing up, and restoring files and directories on vSphere Virtual Machine File System (VMFS).

  The corresponding log file for this component is **/usr/local/avamarclient/var-proxy-N/<Jobname>-<EPOCH>-vmimagew.log**.

- **AvTar**: The primary process for backups and restores. AvTar communicates with GSAN.

  The corresponding log file for this component is **/usr/local/avamarclient/var-proxy-N/<Jobname>-<EPOCH>-vmimagew_avtar.log**.

- **GSAN**: Global Storage Area Network (GSAN) is a component of the VDP appliance. It is sometimes referred to as the data server or the storage server.

  The corresponding log file for this component is **/data01/cur/gsan.log**.

# VDP Installation and Configuration 2

This chapter includes the following topics:


- "vSphere Data Protection Capacity Requirements" on page 22
- "Software Requirements" on page 22
- "System Requirements" on page 23
- "Preinstallation Configuration" on page 23
- "VDP Best Practices" on page 26
- "VDP Installation" on page 29
- "Initial Configuration" on page 30

# vSphere Data Protection Capacity Requirements

vSphere Data Protection (VDP) capacity requirements depend on a number of factors including:

- Number of protected virtual machines
- Amount of data contained in each protected virtual machine
- Types of data being backed up (OS files, documents, and databases, for example)
- Backup data retention period (daily, weekly, monthly, or yearly)
- Data change rates

**NOTE** Assuming average virtual machine sizes, data types, data change rates, and a retention policy of 30 days, 1 TB of VDP backup data capacity supports approximately 25 virtual machines.

# Software Requirements

VDP 6.1 requires the following software:

- vCenter server 5.5 or later

  vCenter server 6.0 or later for VDP 6.1.6 or later

  VDP 6.1 supports the Linux-based vCenter server virtual appliance and the Windows-based vCenter server.

- vSphere web client

  The following link provides information about current vSphere web browser support:

  https://pubs.vmware.com/vsphere-60/index.jsp?topic=%2Fcom.vmware.vsphere.install.doc%2FGUID-F6D456D7-C559-439D-8F34-4FCF533B7B42.html&resultof=%22web%22%20%22client%22

  Web browsers must be enabled with Adobe Flash Player 11.3 or later to access the vSphere web client and VDP functionality.

- vSphere host 5.1 or later

## Hardware Versions and Migration

The virtual machine's hardware version limits virtual machines from migrating to the older versions that are configured on the newer versions of vSphere hosts. If the VDP appliance was migrated to a vSphere host that was version 5.1 or earlier, it would not be functional.

## vSphere Hosts and vSphere Flash Read Cache Compatibility and Performance

The VDP appliance that is deployed as a virtual machine with hardware version 7 enables backward compatibility with vSphere 4.x hosts. The vSphere Flash Read Cache-backed disks are only available on vSphere 5.x and later hosts, which expect a VM to have hardware version 10 or later. As a result, if you attempt to perform an image level backup of a vSphere Flash Read Cache-backed disk by using the VDP appliance, then the current configuration causes the appliance to use the network block device (NBD) protocol (instead of HotAdd) as the transport mode, which adversely affects performance.

## Unsupported Disk Types

- When planning for backups, make sure the disks are supported by VDP. Currently, VDP does not support the following virtual hardware disk types:
  - Independent
  - RDM Independent - Virtual Compatibility Mode
  - RDM Physical Compatibility Mode

### Unsupported Virtual Volumes

The VDP appliance version 6.1 does not support backups and restores of virtual machines on Virtual Volumes (VVOLs).

## System Requirements

The following section lists the system requirements for VDP.

### VDP System Requirements

VDP is available in the following configurations:

- 0.5 TB
- 1 TB
- 2 TB
- 4 TB
- 6 TB
- 8 TB

**IMPORTANT** After VDP is deployed the size can be increased.

VDP requires the following minimum system requirements:

**Table 2-2.** Minimum system requirements for VDP

| | 0.5 TB | 1 TB | 2 TB | 4 TB | 6 TB | 8 TB |
|---|---|---|---|---|---|---|
| Processors | Minimum four 2 GHz processors | Minimum four 2 GHz processors | Minimum four 2 GHz processors | Minimum four 2 GHz processors | Minimum four 2 GHz processors | Minimum four 2 GHz processors |
| Memory | 4 GB | 4 GB | 4 GB | 8 GB | 10 GB | 12 GB |
| Disk space | 873 GB | 1,600 GB | 3 TB | 6 TB | 9 TB | 12 TB |

### IPv6 Requirements

DNS servers that VDP uses in an IPv6 environment must only contain AAAA records for hostnames. The DNS server must not contain both an A and an AAAA record with the same hostname.

## Preinstallation Configuration

Before the VDP installation, complete the following preinstallation steps:

- "DNS Configuration" on page 23
- "NTP Configuration" on page 24
- "User Account Configuration" on page 24
- "VDP Best Practices" on page 26

### DNS Configuration

The DNS server must support both forward and reverse lookup on the VDP and the vCenter.

Before you deploy VDP, you must add an entry to the DNS server for the VDP appliance's IP address and Fully Qualified Domain Names (FQDN). In addition, communication to DNS is required by VMware proxy nodes (port 53) over both TCP and UDP protocols. Failure to set up DNS properly can cause many runtime or configuration issues.

To confirm that DNS is configured properly, run the following commands from the command prompt on the vCenter server:

- `nslookup` <*FQDN_of_VDP*>

  The `nslookup` command returns the FQDN of the VDP appliance.

- `nslookup` <*FQDN_of_vCenter*>

  The `nslookup` command returns the FQDN of the vCenter server.

If the `nslookup` commands return the proper information, close the command prompt. If the `nslookup` commands do not return the information you seek, you can manually add the VDP name and address to the `/etc/hosts` file in the vCenter.

## NTP Configuration

VDP leverages VMware Tools to synchronize time through NTP. All vSphere hosts and the vCenter server must have NTP configured properly. The VDP appliance gets the correct time through vSphere and must not be configured with NTP.

**CAUTION** Configuring NTP directly on the VDP appliance causes time synchronization errors.

See the ESXi and vCenter server documentation for more information about configuring NTP.

## vCenter Hosts and Clusters View

The VDP appliance can work with folders and resource views that are created under the Hosts and Clusters view. The Hosts and Clusters view in the vSphere web client enables you to perform the following tasks:

- Configure user accounts
- Create a snapshot
- Mount the ISO image
- Remove a snapshot
- Revert back to a snapshot
- Expand disks
- Configure the VDP appliance system settings.
- Remove the VDP appliance from the vCenter inventory.

### Accessing the Host and Clusters view

1. From a web browser, access the vSphere web client:

   **https://**<*IP_address_vCenter_Server*>**:9443/vsphere-client/**

2. Log in with administrative privileges.
3. Select **vCenter** > **Hosts and Clusters**.

## User Account Configuration

Before the vCenter user account can be used with VDP, or before the SSO admin user can be used with VDP, you must add these users as administrator on the vCenter root node. Users who inherit permissions from group roles are not valid.

**NOTE** In high-security environments, you can restrict the vCenter user account permissions required to configure and administer the VDP appliance. The account permission categories are listed in "Minimum Required vCenter User Account Permissions" on page 195.

The following steps are used to configure the VDP user or SSO admin user using the vSphere web client.

1 From a web browser, access the vSphere web client:

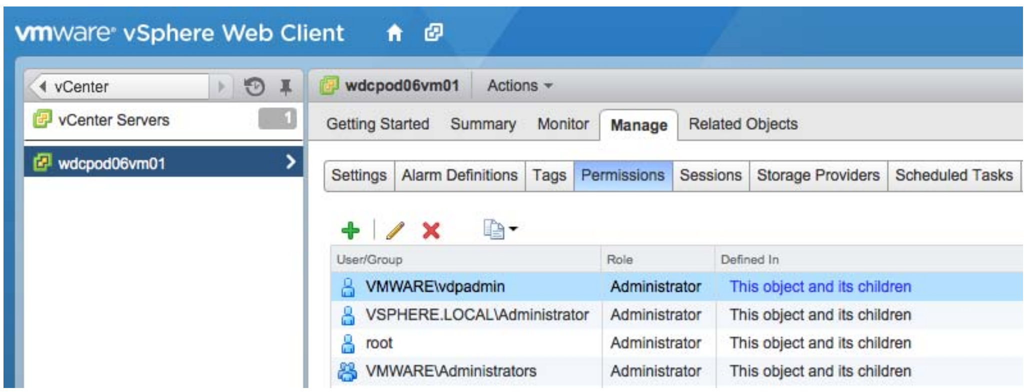
**https://**<*IP_address_vCenter_Server*>**:9443/vsphere-client/**

2 Log in with administrative privileges.

3 Select **vCenter** > **Hosts and Clusters**.

4 On the left side of the page, click the vCenter server.

IMPORTANT Ensure that you select the vCenter from the root level of the tree structure (represented under Hosts and Clusters). If you select the vCenter VM, the configuration fails.




5 Click the **Manage** tab, and then select **Permissions**.

6 Click the **Add permission (+)** icon.

7 Click **Add**.

8 From the Domain drop-down list, select domain, server, or VSPHERE.LOCAL.

NOTE: For vCenter versions 5.1 and earlier, the default domain is SYSTEM-DOMAIN.

9 Select the user who will administer VDP or be the SSO admin user, and then click **Add**.

10 Click **OK**.

11 From the **Assigned Role** list, select **Administrator**.

12 Confirm that the **Propagate to child objects** box is selected.

13 Click **OK**.

To verify that user is listed under Administrators, go to **Home > Administration > Role Manager** and click the **Administrator** role. The user you just added should be listed to the right of that role.

IMPORTANT If the VDP backup user using the VDP Configure utility belongs to a domain account, use the "SYSTEM-DOMAIN\admin" format in VDP-configure. If the username is entered in the format "admin@SYSTEM-DOMAIN" format, tasks related to the backup job may not show up on the Recent Running tasks.

IMPORTANT The domain account password cannot contain spaces.

# VDP Best Practices

The following best practices should be used when deploying, using, and monitoring a vSphere Data Protection (VDP) appliance.

## General Best Practices

- Deploy the VDP appliance on shared VMFS5 or later to avoid block size limitations.
- Make sure that all virtual machines are running hardware version 7 or later to support Change Block Tracking (CBT).
- Install VMware Tools on each virtual machine that VDP will back up. VMware Tools add additional backup capability that quiesces certain processes on the guest OS before the backup. VMware Tools are also required for some features used in File Level Restore.
- When configuring the network for the VDP appliance and the vCenter, do not modify network address information by using NAT or other configuration methods (firewall, IDS, or TSNR). When these unsupported methods are deployed as part of the virtual network, some VDP functionality may not work as designed.

## Deployment Best Practices

The following should always be considered a best practice when deploying a VDP appliance:

- Create a DNS record for the VDP appliance before deploying the appliance. Ensure both forward and reverse lookup are enabled in DNS.
- One of the last steps when deploying VDP is the option to run a storage performance analysis. Run this analysis to verify the storage on which VDP is running meets or exceeds performance requirements. The analysis could take from 30 minutes to a few hours to complete.
- Place the VDP appliance on a separate datastore than where the protected VMs will reside.
- When scheduling backup jobs, consider staggering the start times.
- When applications such as, a database or an Exchange Server with high change rates are backed up, interleave (optimize) them with the image level backups of a different VM that has the unstructured data.

  The databases belong to the structured data and the image level backups belong to the unstructured data. The databases generate more unique data that results in lower deduplication ratio, but the image level backups have higher deduplication ratio. Interleaving the image level backups in between two application backups puts less load on the deduplication engine that results in a better performance of backups.

**NOTE** Generally do not perform the image level backups of an application or a VM that has a high change rate because the snapshot processing can consume time, which impacts the performance.

- Consider other processes that may be running. Try not to schedule replication or automatic backup verification when backup jobs are running. If possible, schedule these jobs to run after the backup jobs have completed, before the maintenance window opens.
- An internal proxy must be activated, and is automatically activated, when an emergency restore operation is performed.

## HotAdd Best Practices

The HotAdd transport mechanism is recommended for faster backups and restores and less exposure to network routing, firewall and SSL certificate issues. If you use the network block device (NBD) transport mechanism instead of HotAdd, backup performance will be degraded.

The following mandatory requirements must be met for a disk to be mounted with HotAdd:

- If you are using vSphere Host version 5.0, the host must be licensed for HotAdd. vSphere Host version 5.1 and later include this feature by default.
- The VDP appliance is deployed on a vSphere host that has a path to the storage that contains the virtual disks being backed up.
- HotAdd is not used on IDE-configured virtual disks. I/O over the network negatively impacts performance. Use SCSI virtual disks instead.
- The total capacity of the VMFS volume, on which VDP resides, is equal to the size of the largest virtual disk being backed up (free space can be less than this size).
- The block size of the VMFS volume where VDP resides is the same or larger than the size of the largest virtual disk being backed up.
- The virtual machine being backed up must not have any independent virtual hard disk.
- The virtual machine being backed up is in the same datacenter (vCenter server container object) as the VDP appliance. HotAdd transport cannot cross the datacenter boundary.
- The virtual machines and VMDKs in the vCenter server have the same name as those associated with the virtual machine being backed up.
- HotAdd does not work with virtual machines that use vSphere Flash Read Cache (vFlash).

For more information about HotAdd best practices, refer to the following Knowledge Base article:

http://kb.vmware.com/kb/2048138

## Storage Capacity for Initial VDP Deployment

When a new vSphere Data Protection (VDP) appliance is deployed, the appliance typically fills rapidly for the first few weeks. This is because nearly every client that is backed up contains unique data. VDP deduplication is most effective when other similar clients have been backed up, or the same clients have been backed up at least once.

After the initial backup, the appliance backs up less unique data during subsequent backups. When initial backups are complete and the maximum retention periods are exceeded, it is possible to consider and measure the ability of the system to store about as much new data as it frees each day. This is referred to as achieving steady state capacity utilization. Ideal steady state capacity should be 80%.

## Sizing

vSphere Data Protection sizing helps determine the VDP appliance size and number of appliances required based on:

- Number of and type of VMs (does the VM contain file system or database data?)
- Amount of data
- Retention periods (daily, weekly, monthly, yearly)
- Typical change rate

The following table shows examples for vSphere Data Protection sizing recommendations:

**Table 2-3.** Sample recommendations for vSphere Data Protection sizing

| # of VMs | Data storage per client | Retention: daily | Retention: weekly | Retention: monthly | Retention: yearly | Recommendation |
|---|---|---|---|---|---|---|
| 25 | 20 | 30 | 0 | 0 | 0 | 1-0.5 TB |
| 25 | 20 | 30 | 4 | 12 | 7 | 1 -2 TB |
| 25 | 40 | 30 | 4 | 12 | 7 | 2 - 2 TB |
| 50 | 20 | 30 | 0 | 0 | 0 | 1 - 1 TB |
| 50 | 20 | 30 | 4 | 12 | 7 | 2 - 2 TB |
| 50 | 40 | 30 | 4 | 12 | 7 | 3 -2 TB |
| 100 | 20 | 30 | 0 | 0 | 0 | 1 - 2 TB |
| 100 | 20 | 30 | 4 | 12 | 7 | 3 - 2 TB |
| 100 | 40 | 30 | 4 | 12 | 7 | 6 - 2 TB |

The recommendations above (note that these are only guidelines) are based on the following assumptions:

- The VMs primarily contain file system data. If the VMs primarily contain database data, the deduplication rates will be lower.
- 70% initial deduplication rate for file system data.
- 99.7% daily deduplication rate for file system data.
- The annual growth rate is 5%.

## Monitoring VDP Capacity

You should proactively monitor VDP capacity. You can view VDP capacity through the **VDP Reports** tab, Used Capacity (which is used to determine steady state). Refer to "Viewing Information from the Reports Tab" on page 108 for more information.

Table 2-4 describes VDP behavior for key capacity thresholds:

**Table 2-4.** Capacity thresholds

| Threshold | Value | Behavior |
|---|---|---|
| Capacity warning | 80% | VDP issues a warning event. |
| Capacity error | 95% | Tasks are not generated on vCenter for backup jobs when capacity is greater than 95% full. |
| Healthcheck limit | 95% | Existing backups are allowed to complete but new backup activities are suspended. VDP issues warning events. |
| Server read-only limit | 100% | VDP transitions to the read-only mode by changing the appliance status to Admin, and does not allow new data. |

Once you exceed 80% capacity, use the following guidelines for capacity management:

- Stop adding new virtual machines as backup clients.
- Remove unnecessary restore points.
- Delete jobs that are no longer needed.
- Reassess retention policies to see if you can decrease retention policies
- Consider adding additional VDP appliances and balance backup jobs between multiple appliances

# VDP Installation

The vSphere Data Protection (VDP) installation is completed through two steps:

- "Deploying the OVF Template" on page 29
- "Initial Configuration" on page 30

## Deploying the OVF Template

### Prerequisites

- vSphere host 5.1 or later.
- vCenter server 5.5 or later.

  vCenter server 6.0 or later for VDP 6.1.6 or later.
- Log in to the vCenter server from a vSphere web client to deploy the OVF template. If you are unable to connect to the vSphere web client, confirm that the vSphere web client service is started.
- The VDP appliance connects to a vSphere host using port 902. If there is a firewall between the VDP appliance and the vSphere Host, port 902 must be open. See Chapter A, "VDP Port Usage," on page 191, for additional information on port usage.
- The VMware Client Integration Plug-in must be installed on your browser. If it is not already installed, it can be installed during the following procedure.

### Procedure

1. From a web browser, access the vSphere web client:

   **https://**<*IP_address_vCenter_Server*>**:9443/vsphere-client/**
2. Log in with administrative privileges.
3. Select **vCenter > Datacenters.**
4. On the **Objects** tab, click **Actions > Deploy OVF Template.**
5. If prompted, allow and install the VMware Client Integration Plug-in.
6. Select the source where the VDP appliance is located. By default the File name dialog is set to OVF Packages (*.ovf). From the drop-down box to the right of File name, select **OVA Packages (*.ova)**.
7. Navigate to the location of the VDP appliance .ova file. Confirm that you select the appropriate file for the datastore. Click **Open**.
8. After the VDP appliance .ova file is selected, click **Next**.
9. Review the template details and click **Next**.
10. On the Accept EULAs screen, read the license agreement, click **Accept**, and then click **Next**.
11. On the Select name and folder screen, type the name for the VDP appliance. When typing the name, use the fully qualified domain name (FQDN), which the VDP configuration uses to find the VDP appliance in the vCenter inventory. Do not change the VDP appliance name after installation.
12. Click the folder or datacenter where you want to deploy the VDP appliance, and then click **Next.**
13. On the Select a resource screen, select the host for the VDP appliance and click **Next**.
14. On the Select Storage screen, select the virtual disk format and select the location of the storage for the VDP appliance. Click **Next**.
15. On the Setup networks screen, select the Destination Network for the VDP appliance and click **Next**.

16 In the Customize template screen, specify the **Default Gateway**, **DNS**, **Network 1 IP Address**, and **Network 1 Netmask**. Confirm that the IP addresses are correct and match the entry in the DNS server. Setting incorrect IP addresses in this dialog box will require the .ova to be redeployed. Click **Next**.

NOTE The VDP appliance does not support DHCP. A static IP address is required.

17 On the Ready to complete screen, confirm that all of the deployment options are correct. Check **Power on** after deployment and click **Finish**.

vCenter deploys the VDP appliance and boots into the install mode. You can monitor **Recent Tasks** to determine when the deployment completes.

# Initial Configuration

## Prerequisites

- Ensure that the VDP `.ovf` template was deployed successfully. See "Deploying the OVF Template" on page 29 for more information.
- You must be logged into the vCenter server from the vSphere web client.
- Enough free disk space exists on the datastore. When an optional performance analysis test is run during initial configuration of the appliance, 41 GB per disk per datastore is required (for example, if three disks are placed on the same datastore, 123 GB of free space is required). If there is not enough space available, the test reports a value of 0 for all of the read, write, and seek (optional) tests and gives a final status of insufficient space.

## Procedure

1 From a web browser, access the vSphere web client:

**https://**<*IP_address_vCenter_Server*>**:9443/vsphere-client/**

2 Log in with administrative privileges.

3 Select **vCenter Home > vCenter > VMs and Templates**. Expand the vCenter tree and select the VDP appliance.

4 Open a console session into the VDP appliance by right-clicking the VDP appliance and selecting **Open Console**.

5 After the installation files load, the Welcome page for the VDP menu appears. Open a web browser and type:

**https://**<IP_address_VDP_appliance>**:8543/vdp-configure/**

The VDP Login page appears.

6 Type **root** in the **User** field and **changeme** in the **Password** field, and then click **Login**.

The VDP Welcome page appears.

7 Click **Next**.

The Network Settings dialog box appears by default.

8 Either confirm or specify the following network and server information for your VDP appliance. Ensure that the values are correctly populated. Otherwise, the installation fails.

a IPv4 or IPv6 static address

b Netmask

c Gateway

d Primary DNS

e Secondary DNS

f Hostname

g Domain

9 Click **Next**.

The Time Zone dialog box appears.

10 Select the appropriate time zone for your VDP appliance, and click **Next**.

The VDP Credentials dialog box appears.

11 In the **Password** field, type in the VDP appliance password by using the following criteria, and then verify the password by retyping it in the **Verify password** field. This password is the universal configuration password.

The four-character classes are as follows:

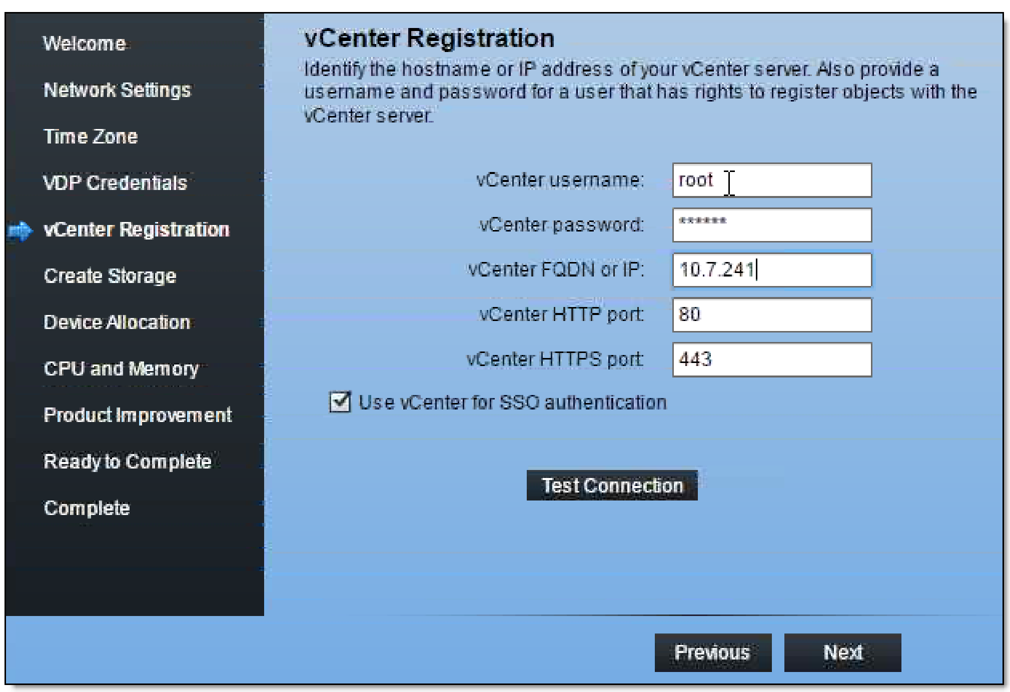
- Upper case letters A-Z
- Lower case letters a-z
- Numbers 0-9
- Special characters (for example: ~!@#,.)

Create the password using the following criteria:

- If all four character classes are used, the password must be at least 6 characters.
- If three character classes are used, the password must be at least 7 characters.
- If one or two character classes are used, the password must be at least 8 characters.

12 Click **Next**.

The **vCenter Registration** page appears.



13 Type in the following fields:

- vCenter username

  If the user belongs to a domain account, enter the name by using the "SYSTEM-DOMAIN\admin" format.

**CAUTION** If an SSO admin user is specified as the vCenter username in the format <username@vsphere.local>, tasks related to VDP operations do not appear in the vCenter Recent Tasks pane of the vSphere web client. For tasks to appear in the Recent Tasks pane, specify the SSO admin user in the format <vsphere.local\username>.

- vCenter password
- vCenter FQDN or IP
- vCenter HTTP port

  The default port is **80**. If you use a port other than the default port, you must open the port in **/etc/firewall.base** and restart the **avfirewall** service.

  Specify a custom value for the HTTP port if you need to connect to vCenter over the HTTP port, instead of the HTTPS port, which is used for all other communication.

- vCenter HTTPS port

  The default port is **443**. If you use a port other than the default port, you must open the port in **/etc/firewall.base** and restart the **avfirewall** service.

  If the field is disabled, select the **Use vCenter for SSO authentication** checkbox for SSO authentication.

**NOTE** Leave the **Use vCenter for SSO authentication** checkbox enabled if your vCenter has SSO embedded in the vCenter server appliance. If you disable this selection by clearing the checkbox, you must enter the SSO Server FQDN or IP address and the SSO port fields.

- Click **Test Connection**.

  A connection success message appears. If this message does not appear, troubleshoot your settings and repeat this step until a successful message appears.

If you receive the following message on the vCenter Registration page, perform the steps that "User Account Configuration" on page 24 describes:

```
Specified user either is not a dedicated VDP user or does not have sufficient vCenter
privileges to administer VDP. Please update your user role and try again.
```

14 Click **Next** to advance to the **Create Storage** page, which guides you through the storage type selection. "Creating New Storage" on page 72 provides information about storage configuration, and the final steps that you require to complete the initial configuration wizard.

# VDP Appliance Migration

3

This chapter contains the following topic:


- "Migrating the VDP Appliance from 5.8.x or Later to 6.1" on page 36

## Migrating the VDP Appliance from 5.8.x or Later to 6.1

VDP 6.1 does not support upgrading from previous versions. However, you can migrate all jobs, maintenance window, and data from the VDP 5.8.x or later appliance to the VDP 6.1 appliance.

**NOTE** Do not perform migration on a Replication Target Appliance that is a VDP appliance, on which no jobs have been created, but used as a replication target for data. To upgrade a replication target, perform the steps that "Attaching Existing VDP Disks" on page 74 describes.

1. Start the VDP initial configuration wizard.
2. Specify the required fields on the pages **Welcome** through **vCenter Registration**.
3. On the **Create Storage** page:
   a. Select **VDP Migration**.
   b. Specify the following fields:
      - **IP Address**: Type the IP address of the VDP 5.8.x or later appliance.
      - **Password**: Type the password of the VDP 5.8.x or later appliance.
   c. Click **Verify authentication**.

      Ensure that the authentication succeeds.
   d. Click **Next**.
   e. In the **Warning** message box which states `Migrating VDP will shut down your source VDP <xx.xx.xx.xx>. Do you wish to continue?`, click **Yes**.
   f. Click **Next**.
4. On the **VDP Migration** page:
   a. Ensure that the progress and the log messages properly appear.
   b. Click **Next**.
5. On the **Device Allocation** page:
   a. Click the browse button to browse the ESX datastore, and select the data disks of VDP in any order to migrate.
   b. Ensure that the disk validation succeeds.
   c. Click **Next**.
6. Specify the required fields on the pages **CPU and Memory** through **Complete**.

# Post-Installation Configuration of VDP Appliance

4

This chapter contains the following topics:


- "About the VDP Configure Utility" on page 38
- "Viewing Status" on page 38
- "Starting and Stopping Services" on page 39
- "Collecting VDP Logs or Diagnostic Information" on page 40
- "Modifying Configuration Settings" on page 42
- "Rolling Back an Appliance" on page 43
- "Emergency Restore" on page 44

# About the VDP Configure Utility

During installation of vSphere Data Protection (VDP), the VDP Configure utility runs in “install” mode. This mode allows you to enter initial networking settings, time zone, VDP appliance password, and vCenter credentials. Install mode also allows you to create or attach storage and optionally run the performance assessment tool. After initial installation, the VDP Configure utility runs in “maintenance” mode and displays a different user interface.

To access the VDP Configure utility, open a web browser and type:

**https://**<*IP_address_VDP_Appliance*>**:8543/vdp-configure/**

Use the VDP appliance username and password. As you log into the VDP configuration utility, a system health check script runs. You must wait for the system health check to finish running before you can perform configuration tasks from any of the VDP configuration utility tabs.

**NOTE** After the appliance is configured, you can optionally run a performance analysis test from the VDP configuration utility. The test requires 41 GB per disk per datastore. If that amount of space is not available, the utility reports insufficient space and the performance analysis test will not run.

Use the configuration interface to perform the following configuration tasks:

- “Storage Management” on page 71: Enables you to view your storage configuration, add and edit a Data Domain system, and add or expand disk storage on the VDP appliance.
- “Viewing Status” on page 38: Enables you to see the services currently running (or currently stopped) on the VDP appliance.
- “Starting and Stopping Services” on page 39: Enables you to start and stop selected services on the VDP appliance.
- “Collecting VDP Logs or Diagnostic Information” on page 40: Enables you to download current logs from the VDP appliance for troubleshooting purposes.
- “Modifying Configuration Settings” on page 42: Enables you to view or change network settings, configure vCenter Registration, view or edit system settings (time zone information and VDP credentials), and manage proxy throughput through external proxy and internal proxy configuration options.
- “Rolling Back an Appliance” on page 43: Enables you to restore the VDP appliance to an earlier known and valid state.
- “Emergency Restore” on page 44: Enables you to restore a VM directly to the host that is running the VDP appliance. This emergency restore procedure is intended for use when the vCenter is unavailable.
- “Managing Proxy Throughput” on page 65: Enables you to manage internal and external proxies. You can deploy up to 8 external proxy virtual machines. Deploying 8 external proxies helps you to improve the backup throughput by using hot add feature, and run more concurrent backup and restore operations.

# Viewing Status

The **Configuration** tab lists all of the services required by VDP and the current status of each service. Table 4-5 describes the services used by VDP.

**Table 4-5.** Description of services running on the VDP appliance.

| Service | Description |
|---|---|
| Core services | These are the services that comprise the backup engine of the VDP appliance. If these services are disabled no backup jobs (either scheduled or “on demand”) will run, and no restore activities can be initiated. |
| Management services | Management services should only be stopped under the direction of technical support. |
| File level restore services | These are the services that support the management of file-level restore operations. |
| Backup recovery services | These are the services that support backup recovery. |

**Table 4-5.** Description of services running on the VDP appliance. (Continued)

| Service | Description |
|---|---|
| Maintenance services | These are the services that perform maintenance tasks, such as evaluating whether retention periods of backups have expired. The Maintenance services are disabled the first 24-48 hours after the VDP appliance is deployed. This creates a larger backup window for initial backups. |
| Backup scheduler | The backup scheduler is the service that initiates schedule backup jobs. If the backup scheduler is stopped, no scheduled backups will run. On-demand backups can still be initiated. |
| Replication service | This service replicates backups of one VDP to either another VDP or Avamar. |

**NOTE** If any of these services stop running, an alarm is triggered on the vCenter server. If a stopped service is restarted, the alarm is cleared. There can be a delay of up to 10 minutes before alarms are triggered or cleared.

The status that is displayed for these services can be any of the following:

- Starting
- Start Failed
- Running
- Stopping
- Stop Failed
- Stopped
- Loading-getting state
- Unrecoverable (Core services only)
- Restoring (Management services only)
- Restore Failed (Management services only)

Click the refresh icon to update the status display.

## Starting and Stopping Services

On the status screen you can restart stopped services by clicking **Start**, or you can stop running services by clicking **Stop**. In general, however, you should only stop running services under the direction of technical support.

If you see that a service is stopped, you can attempt to restart it by clicking **Start.** In some cases, however, additional troubleshooting steps are necessary for the service to work properly.

If all services are stopped, start the services in the following order:

1. Core services
2. Management services
3. Backup scheduler
4. Maintenance services
5. File level restore services
6. Backup recovery services
7. Replication service

# Collecting VDP Logs or Diagnostic Information

The **Log Collector** tab enables you to individually or collectively download the VDP log files. The **Log Collector** tab groups log files into four sections. Table 4-6 describes the log files in each group.

**Table 4-6.** Log files available from the Log Collector tab

| Group | Log files |
|---|---|
| All VDP Appliance Logs | Log files related to the VDP appliance |
| Client Logs | Log files related to Microsoft applications |
| Configurations | VDP configuration files from proxies, config checker, agents, and so on<br>These configuration files are located in `/space/vdp/config`. |
| External Proxy Logs | Log files from external proxy virtual machine<br>These log files are visible only when at least one external proxy is deployed. |

**NOTE** All logs are located in `/space/vdp/logs`.

The All VDP Appliance Logs group is divided into following subgoups so that you can download only specific log files:

**Table 4-7.** Log files available in the All VDP Appliance Logs group

| Group | Log files |
|---|---|
| Core VDP Service | Log files related to GSAN, VDP, AVI, system, and DPN |
| Management Service | Log files related to the MC server |
| File System Service | Log files related to HFScheck |
| File Level Restore Service | Log files related to FLR |
| Replication | Log files related to replication and replication recovery |
| Image Backup and Restore | Log files related to backups and restores |

### Procedure

1 Open a web browser and type:

**https://**<*IP_address_VDP_Appliance*>**:8543/vdp-configure/**

2 Log in with the VDP username and password.

3 On the **Log Collector** tab, select one or more options and click **Download** to download log files:

- **All VDP Appliance Logs**: Select this option and click **Download** to download all log files from the VDP services to a `.zip` file.

  The **Select location for download** dialog box appears. The default name of the log bundle is `LogBundle.zip`.

  Rename the file to a unique name. The log bundle is intended primarily to send VDP appliance logs to the Support personnel.

  - To download all logs under a specific heading from the following list, select a checkbox next to the heading and click **Download**.
  - To download the log files that are listed under multiple headings, select the checkbox next to one or more log files and click **Download**.

- **Client Logs**: Select this option and click **Download** to download an aggregated text file that contains all client failure logs.
- **Configurations**: Select this option and click **Download** to download only the VDP configuration file information.

The following table provides information about what log files you must download to troubleshoot the area-specific issues:

**Table 4-8.** Log files to download to troubleshoot the area-specific issues

| Area to troubleshoot | Log group | Location on appliance | Relevant log files to download |
|---|---|---|---|
| Upgrade | Core VDP Service -> AVI | /space/vdp/logs/avi_logs/server_log/ | avinstaller.log.0 |
| Image Backup or Restore | Image Backup and Restore and External proxy Logs (if you are using external proxy) | ■ /space/vdp/logs/image_proxy/avamarclient/<br>■ /usr/local/avamarclient/var (in case of external proxy) | ■ *.log files (also known as avtar logs and contain information on the backup)<br>■ *.alg files (contains workorder information) |
| Replication | Replication | /space/vdp/logs/replicate/client/ | ■ *Replicate-avreplscript.log<br>■ *Replicate.log<br>■ *Replicate.alg |
| File Level Recovery | File Level Restore Service | ■ /usr/local/avamarclient/bin/logs<br>■ /space/vdp/logs/flr_proxy/<br>■ /space/vdp/logs/vdp_logs/flr/server_logs/ | ■ flr-server.log<br>■ *.log<br>■ flr_debug.txt<br>■ flr_msg.txt<br>■ flr_out.txt |
| VDP plug-in | Core VDP Service -> VDP | /space/vdp/logs/vdp_logs/vdr/server_logs/ | vdr_server.log |
| VDP configuration utility | Core VDP Service -> VDP | /space/vdp/logs/vdp_logs/vdr/server_logs/ | vdr-configure.log |
| Integrity check | File System Service -> HFSCheck | /space/vdp/logs/hfscheck/ | *.log |

The following table provides information about what logs you must collect to troubleshoot the backup and restore issues with Microsoft applications:

**Table 4-9.** Log files to download to troubleshoot the area-specific issues

| Plug-in | Location of log files | Relevant log files to collect |
|---|---|---|
| SQL VDP plug-in | ■ For SQL standalone:<br>C:\Program Files\avp\var<br>■ For SQL failover cluster:<br>The var folder on a share drive that was specified during the cluster configuration by using Windows Cluster Configuration Wizard (WCCW)<br>■ For SQL AlwaysOn cluster:<br>The var folder that was specified during the cluster configuration by using WCCW<br>If a shared drive was specified during the cluster configuration, the shared drive contains the log files of all the nodes.<br>If a local folder was specified during the cluster configuration, each cluster node contains the log files. | ■ SQL VSS plug-in's logs<br>■ Avtar logs<br>■ Avagent logs<br>■ Windows Application and System event viewer logs on the client |
| SharePoint VDP plug-in | C:\Program Files\avp\var | ■ SharePoint VSS plug-in's logs<br>■ Avtar logs<br>■ Avagent logs<br>■ Windows Application and System event viewer logs on the client |

**Table 4-9.** Log files to download to troubleshoot the area-specific issues

| Plug-in | Location of log files | Relevant log files to collect |
|---|---|---|
| NOTE<br>■ In a SharePoint federated farm, you must collect the logs from every SharePoint Server node including any SQL backends.<br>■ If the SharePoint Server is connected to either a SQL failover or an AlwaysOn cluster, or both a SQL failover and an AlwaysOn cluster, a set of logs is also generated on the shared var folder that was specified during the cluster configuration by using WCCW. | | |
| Exchange VDP plug-in | ■ On each VDP client:<br>C:\Program Files\avp\var<br>■ Shared var folder that was spefified in the Windows Cluster Configuration Wizard for Exchange DAG | ■ Exchange plug-in's logs<br>■ Avtar logs<br>■ Avagent logs<br>■ Windows Application and System event viewer logs on the client |

# Modifying Configuration Settings

When you access the VDP Configure utility after installation, it runs in maintenance mode. While in maintenance mode, by clicking the **Configuration** tab, you can set or modify any settings that were entered during installation. You can configure network settings and system settings, and you can register the vCenter.

## Network Settings

You can configure the following network settings on the **Configuration** tab.

NOTE If you change any network settings, manually run an integrity check immediately after the change is made. Failure to run an integrity check could result in a VDP connection failure while rolling back to a previous checkpoint. See "Running an Integrity Check" on page 59 for instructions.

- **IPv4 or IPv6 static address**: The Internet Protocol v4 or v6 setting of the interface.
- **Netmask/Prefix**: The network mask or prefix of the IPv4 or IPv6 static address.
- **Gateway**: The gateway address of the IPv4 or IPv6 static address.
- **Primary DNS**: The primary domain name system used for DNS resolution.
- **Secondary DNS**: The secondary domain name system used for DNS resolution.
- **Hostname**: The unique name by which a computer is known on a network (for example: vdp-primary).
- **Domain**: A unique name that identifies a website (for example: emc.com). *Domain* is sometimes referred to as the *domain name*.

## vCenter Server Registration

CAUTION Changing the vCenter server hostname, vCenter server password, vCenter server FQDN or IP address, or vCenter server port number results in the loss of all backup jobs, replication jobs, and backup verification jobs associated with the VDP appliance. Existing restore points remain intact, but you will be required to re-create backup jobs and replication jobs.

Before the process that reconfigures the vCenter server registration runs, all of the following conditions must be true:

- No other reconfiguration process is running (for example, password or network reconfiguration).
- No backup, replication, or restore jobs are running.
- No integrity check is running.
- All services are running on the VDP appliance.

If you change any vCenter registration credentials, manually run an integrity check immediately after the change is made. Failure to run an integrity check could result in a VDP connection failure while rolling back to a previous checkpoint. See "Running an Integrity Check" on page 59 for instructions.

## Rolling Back an Appliance

The VDP appliance could become inconsistent or unstable. In some cases, the VDP Configure utility can detect an inconsistent or unstable state and will provide a message similar to the following message immediately after you log in:

> It appears that your VDP appliance has suffered an unclean shutdown and will likely require a checkpoint rollback to restore data protection functionality. This process may be initiated from the 'Rollback' tab.

**CAUTION** By default, VDP keeps two system checkpoints. If you roll back to a checkpoint, all backups and configuration changes taken since the checkpoint was taken will be lost when the rollback is completed.

The first checkpoint is created when VDP is installed. Subsequent checkpoints are created by the Maintenance service. This service is disabled for the first 24 to 48 hours of VDP operation. In the event that you roll back during this time frame, then the VDP appliance is set to the default configuration and any backup configurations or backups are lost.

**NOTE** If any VMware VDP for Exchange Server Clients or VMware VDP for SQL Server Clients were installed between a checkpoint and a rollback, the clients must be reinstalled.

Follow the procedure below to roll back a VDP appliance.

**CAUTION** Only roll back to the most recent validated checkpoint.

### Prerequisite

The VDP appliance must be installed and the VDP appliance password is required.

### Procedure

1 Open a web browser and type:

**https://**<*IP_address_VDP_Appliance*>**:8543/vdp-configure/**

2 Log in with the VDP username and password.

3 Click the **Rollback** tab.

4 Click the lock icon to enable VDP rollback.

5 Enter the VDP appliance password and click **OK**.

The lock icon changes to unlocked.

6 Click the checkpoint that you want to roll back to.

7 Click **Perform VDP rollback to selected checkpoint**.

A warning message appears explaining the consequences of rolling back the VDP appliance.

8 Click **Yes**.

An information message appears telling you a rollback has been initiated.

9 Click **OK**.

The VDP appliance attempts to roll back and displays status information. It also displays an information message indicating whether the roll back succeeded or failed.

10 Click **OK**.

If the VDP appliance did not roll back successfully, contact Customer Support.

# Emergency Restore

VDP is dependent on the vCenter server for many of its core operations. The direct-to-host emergency restore operation provides a method to restore the virtual machines when wither the vCenter server is unavailable or the user is unable to access the VDP user interface by using the vSphere web client.

An emergency restore operation restores a VM directly to the host that is running the VDP appliance. The **Emergency Restore** tab displays a list of VMs that have been backed up by the VDP appliance. These virtual machines can be restored as new virtual machines on to the host where the VDP appliance is running.

## Best Practices and Recommendations

Ensure the following requirements are met before you perform an emergency restore operation:

- The virtual machine being restored has a virtual hardware version that is supported by the host on which the VDP appliance is running.
- There is adequate free space in the target datastore to accommodate the entire virtual machine.
- The target VMFS datastore to which the virtual machine is being restored supports the VMDK file size.
- Network connectivity is available for the restored virtual machines from the host where the VDP appliance is running.
- There is at least one local account with administrator privileges on the host where the VDP appliance is running.

## Limitations and Unsupported Features

- The vSphere host on which the emergency restore operation is being performed cannot be part of the vCenter inventory. A vSphere host that is currently managed by the vCenter server must be temporarily disassociated from the vCenter server to perform the emergency restore. To disassociate the vCenter server, use the vSphere Client (not the vSphere web client) connected directly to the vSphere host.

  In case of ESXi 6.5, use a web browser to connect to the ESXi 6.5 host, and then use **Actions** > **Disconnect from vCenter Server** option on the vCenter Server.
- Emergency restore allows you to restore only to the root of the host level in the inventory.
- Emergency restore requires that the DNS server used by VDP is available and can fully resolve the target vSphere host name.
- Emergency restore restores the virtual machine in the Powered Off state. You must manually log in to the host and power on the restored virtual machine.
- Emergency restore restores the virtual machine as a new virtual machine. You must ensure that the name provided for the virtual machine is not a duplicate of a virtual machine that already exists.
- Emergency restore does not list MSapp clients.
- An internal proxy is automatically activated when an emergency restore operation is performed. If both the internal and external proxies are activated, you must disable the internal proxy on the VDP Configure utility for the emergency restore to complete successfully.

## Procedure

1. If you have not already done so, log in to the vSphere client of the host and, from the **Summary** tab under Host Management, perform the following steps:
   a. Click **Disassociate Host from vCenter**.
   b. Click **Yes** when prompted to remove the host from the vCenter.
2. Log in to the VDP Configure utility:

   **https://**<*IP_address_VDP_Appliance*>**:8543/vdp-configure/**

3 Click the **Emergency Restore** tab.

Virtual machines protected by VDP are listed in the Emergency Restore dialog box. Here, you can find the following details about the virtual machines:

- **Name**: The name of the virtual machines protected by the VDP appliance. By clicking the disclosure arrows, you can determine the date and time of the last restore for the selected virtual machine.
- **Last Known Path**: The last known location of the virtual machine in the vCenter inventory list. This location is updated if the virtual machine is moved.
- Running restore details:
    - **Client Name**: The name of the virtual machine client that is restored.
    - **Status**: The pass or fail status of the restore.
    - **Start Time**: The time the restore started.
    - **Completed Time**: The time the restore completed.
    - **Bytes Transferred**: The number of bytes that were transferred during restore.

4 Select the virtual machine that will serve as the restore point and click **Restore**.

The Host Credentials dialog box appears.

5 In the Host Credentials dialog box, enter valid host credentials:

- **ESXi Host name or IP**: Enter the vSphere hostname or vSphere host IP address.
- **Port - 443**: Populated by default.
- **Username**: Enter the vSphere host username. The recommended host username is "root." For any other host username, the user account must have the create VM privilege.
- **Password**: Enter the vSphere host password. If you enter invalid host credentials, an error message displays and you will be unable to connect to the host.

**NOTE** If you did not successfully disassociate the selected virtual machine from its vCenter, an error message appears and you cannot proceed.

6 Click **OK**.

The Restore a Backup screen initiates the restore with the new name and destination.

7 The Restore a Backup dialog box displays the following information:

- **Client name**: The name of the client on which the virtual machine is restored.
- **Backup**: The date and timestamp of the backups.
- **New Name**: The field where a new name must be entered, which cannot be a duplicate of a VM that already exists.
- **Destination**: The vSphere hostname.
- **Datastore**: A drop-down list of datastores available as the destination targets.

8 Enter a new name in the **New Name** field. The name must be unique and can be up to 255 characters long. The following characters cannot be used in the name: ~ ! @ $ ^ % { } [ ] | , ` ; # \ / : * ? < > ' " & . In addition, diacritical characters cannot be used (for example: â, é, ì, ü, and ñ).

9 Select a datastore as the destination target for the backup.

**CAUTION** The datastore capacity size is listed. Make sure you select a datastore with enough disk space to accommodate the restore. Insufficient space causes the restore to fail.

10 Click **Restore**.

11 Verify that the restore submitted successfully by checking the progress in the Recent Tasks dialog box.

**NOTE** The restored virtual machine is listed at the vSphere host level in the inventory. Restoring to a more specific inventory path is not supported.

## Automatic Host Detection

For vSphere Data Protection versions 5.5 and earlier, users are required to identify, disassociate, and populate vSphere host values before they perform an emergency restore. With vSphere Data Protection version 5.8 and later, the appliance automatically detects the host that it is currently registered to and pre-populates the hostname or IP value in the Host Credentials dialog box. This eases the burden for a user, who might have numerous hosts in a cluster and needs to identify the current, resident host to disassociate from the vCenter.

The following are rare cases when the appliance fails to detect the most updated host to which the appliance is registered and displays an older value:

- The vCenter is unavailable and, after the VDP appliance migrates to a different host under HA-enabled clustering, the appliance displays the older host to which it was registered.
- The vCenter becomes unavailable immediately after the appliance is migrated to another host. In this case, the new host cannot be detected, because it takes the appliance a period of time to process and update events from the vCenter.

In both of these cases, the user must manually determine to which host the appliance is registered.

## Refreshing Restore Points

1 Log in to the VDP-Configure URL:

   **https://***<IP_Address_VDP_Appliance>***:8543/vdp-configure/**

2 Click the **Emergency Restore** tab.

3 Click **Refresh**.

   The loading bar refreshes the restore points.

## Reconnecting the Host to the vCenter

1 Restore the vCenter. Refer to "Restore Operations" on page 132 for instructions.

**NOTE** The restored vCenter is powered off by default.

2 Once the vCenter restore completes, power on the vCenter.

3 Log in to the vCenter URL to verify all the services are running:

   **https://<***IP_address_vCenter***>:5480**

   For a Windows-based vCenter, log in as an administrator to the server that runs vCenter server, and verify whether all the Windows services are running.

4 Log in to the restored vCenter through the vSphere client:

   **https://***<IP_address_vCenter>***:9443/vsphere-client/**

5 From the vSphere client, add the vSphere host to the newly-restored vCenter.

**NOTE** Once the vCenter has been restored, there may be a delay of approximately 20 minutes while the vCenter services start up. During this delay, you will be unable to perform a successful backup or restore operation. If you experience delays, try the backup or restore later.

# Securing Communication between VDP and vCenter

5

This chapter contains the following topics:


- "Replacing the VDP Certificate" on page 48
- "Authenticating the VDP Server for Secured Communications" on page 49
- "Securing the VDP Appliance Communication with the vCenter Server" on page 50
- "Securing the Proxy Communication with the vCenter Server" on page 51

# Replacing the VDP Certificate

You can replace the VDP certificate for the port 8543 with either a new self-signed certificate or a CA-signed certificate.

Perform the following steps to replace the VDP certificate:

1 Stop the services by running the following command:

```
root@vdp#: emwebapp.sh --stop
```

2 Delete the **tomcat** alias by running the following command:

```
root@vdp#:/usr/java/latest/bin/keytool -delete -alias tomcat
```

3 Generate an SSL certificate by running the following command:

```
root@vdp#:/usr/java/latest/bin/keytool -genkeypair -v -alias tomcat -keyalg
RSA -sigalg SHA256withRSA -keystore /root/.keystore -storepass changeit
-keypass changeit -validity 3650 -dname "CN=vdp.vmware, OU=VDP, O=OrgName,
L=PUNE, S=MH, C=IN"
```

4 To use a self-signed certificate, go to step 6.

5 To use a CA-signed certificate:

a Generate a Certificate Signing Request (CSR) by running the following command:

```
root@vdp#:/usr/java/latest/bin/keytool -certreq -keyalg RSA -alias tomcat
-file csrFileName.csr
```

b Get the CA-signed certificate in the **.p7b** format by using the CSR content and save the certificate.

c Import the .p7b chain certificate file by running the following command:

```
root@vdp#:/usr/java/latest/bin/keytool -import -alias tomcat -keystore
/root/.keystore -file /<Path_to_the_Certificate>/<chainCertFileName.p7b>
```

6 Verify whether the certificate entry with the **tomcat** alias exists in the keystore by running the following command:

```
root@vdp#:/usr/java/latest/bin/keytool -list -v -keystore /root/.keystore
-storepass changeit -keypass changeit
```

7 If the certificate entry exists in the key store, run the **addFingerprint.sh** script:

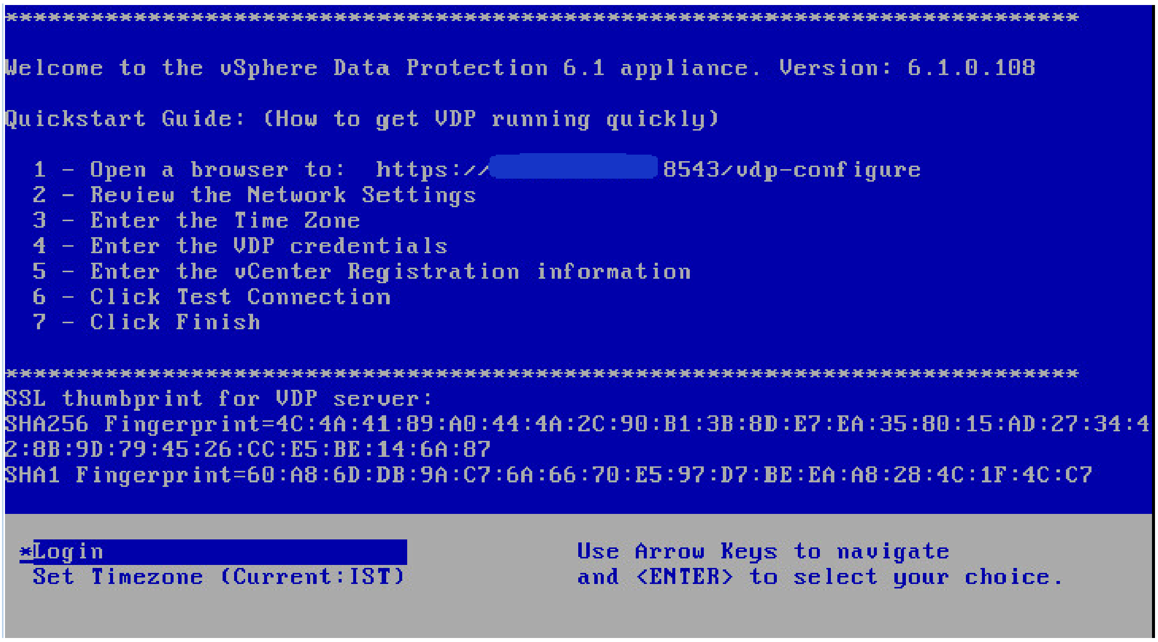
```
root@vdp#:./usr/local/avamar/bin/addFingerprint.sh
```

8 Start the services by running the following command:

```
root@vdp#: emwebapp.sh --start
```

## Authenticating the VDP Server for Secured Communications

While communicating with the VDP server, to prevent the man-in-the-middle attacks and fully use the security that the certificates provide, VDP provides thumbprint on the console. The following figure illustrates a sample thumbprint:




When you access the VDP server by using the **https://<vdp-ip>:8543/vdp-configure** URL, ensure that the SSL certificate comes from the same VDP server, and is valid. To verify whether the SSL certificate is valid, see whether the thumbprint on the console and the SSL certificate that the browser receives during SSL handshake match. VDP has a default SSL certificate.

To change the default certificate, import the new certificate as a **privatekeyentry** in keytool to keystore.

After you import the new certificate, to see the updated thumbprint on the VDP console, perform the following steps:

1 Stop services by running the following command:

   **emwebapp.sh --stop**

2 Run addFingerprint.sh script:

   **./usr/local/avamar/bin/addFingerprint.sh**

3 Start services by running the following command:

   **emwebapp.sh -start**

4 Restart the VDP appliance.

## Securing the VDP Appliance Communication with the vCenter Server

VDP appliance can establish HTTPS connections to vCenter server either with or without SSL certificate verification for different purposes. Failure to verify the SSL certificate makes the vCenter server and the VDP appliance prone to the man-in-the-middle attacks. VDP now supports uploading of the vCenter server certificate or signing certificate in the VDP certificate repositories during the registration process to the vCenter server. The certificate is then used to authenticate or trust vCenter server in all the subsequent HTTPS connections including test connections to the vCenter server, thereby securing the VDP appliance.

### Prerequisites

Download the vCenter certificate in the **.p7b**, **.cer**, or **.crt** format by using the following URL:

**https://**<*vCenter_IP_address*>**/sdk**

### Procedure

- For initial configuration:
  - a Open the initial configuration wizard by using the following URL:

    **https:**<*IP_address_VDP_appliance*>**:8543/vdp-configure/**
  - b On the **vCenter Registration** page:
    - i Specify the vCenter information such as login credentials, host name or IP address, and port numbers in the corresponding fields.
    - ii Select **Verify vCenter certificate**.

      If you do not want to secure the VDP application communication with vCenter, clear **Verify vCenter certificate**. Clearing this option does not verify the vCenter certificate.
    - iii Click **Upload** and select the vCenter certificate that you downloaded.
    - iv Click **Test Connection**.

      If the vCenter certificate is valid, connection to vCenter is established. Otherwise, the connection fails.
- For vCenter re-registration:
  - a Log in to the VDP configuration utility by using the following URL:

    **https:**<*IP_address_VDP_appliance*>**:8543/vdp-configure/**
  - b On the **Configuration** tab, click the setting icon and select **vCenter registration**.

    The **vCenter Registration** wizard appears.
  - c On the **vCenter Registration** page, select **I have reviewed the information. I want to reconfigure vCenter** and click **Next**.
  - d On the **vCenter Configuration** page:
    - i Specify the vCenter information such as login credentials, host name or IP address, and port numbers in the corresponding fields.
    - ii Select **Verify vCenter certificate**.
    - iii Click **Upload** and select the vCenter certificate that you downloaded.
    - iv Click **Next**.
  - e On the **Ready to Complete** page, review the information and click **Finish**.

    If the vCenter certificate and other information is valid, connection to vCenter is established. Otherwise, the re-registration fails.

# Securing the Proxy Communication with the vCenter Server

By default, the VDP server management service and proxies do not validate SSL certificates when connecting to the vCenter server. This can leave the vCenter server vulnerable to a man-in-the-middle exploitation, which might result in unauthorized access to the vCenter server. Configuring each proxy to use SSL certificate authentication when connecting to the vCenter server corrects this vulnerability.

### Prerequisites

Ensure that a Certificate Authority (CA) signed SSL certificate is installed on the vCenter server.

http://kb.vmware.com/kb/2034833 provides information about how to generate a CA-signed SSL certificate and install it on the vCenter server, and implement the CS-signed certificates in a vSphere 5.1 or 5.5 environment.

http://kb.vmware.com/kb/2111219 provides information about how to replace the default certificates with the CA-signed SSL certificates in a vSphere 6.0 environment.

### Procedure

1 Open a command shell and log in to the proxy as `root`.

2 Copy the vCenter server certificate to `/usr/local/avamarclient/bin` on the proxy.

- For a Linux vCenter:
  - vCenter 5.x:
    - Server certificate: `/etc/vmware-vpx/ssl/rui.crt`
    - Root certificate: `/etc/vmware-vpx/ssl/rui-ca-cert.pem`
  - vCenter 6.0:
    - Server certificate: `/etc/vmware-vpx/ssl/rui.crt`
    - Root certificate: `/var/lib/vmware/vmca/root.cer`

  scp the appropriate certificate from vCenter to `/usr/local/avamarclient/bin` on the proxy.

  In case of a self-signed certificate, use only server certificate.

  In case of a CA-signed certificate, create a chain certificate by running the following command:

  cat rui.crt root.cer > chain_cert.pem

- For a Windows vCenter:
  - vCenter 5.x:
    - Server certificate: `C:\ProgramData\VMware\VMware VirtualCenter\SSL\rui.crt`
    - Root certificate: `C:\ProgramData\VMware\VMware VirtualCenter\SSL\cacert.pem`
  - vCenter 6.0:
    - Server certificate: `C:\ProgramData\VMware\vCenterServer\cfg\vmware-vpx\ssl\rui.crt`
    - Root certificate: `C:\ProgramData\VMware\vCenterServer\data\vmca\root.cer`

  scp the appropriate certificate from vCenter to `/usr/local/avamarclient/bin` on the proxy.

  In case of a self-signed certificate, use only server certificate.

  In case of a CA-signed certificate, create a chain certificate by running the following command:

  cat rui.crt root.cer > chain_cert.pem

**NOTE** If a chain SSL certificate is used for vCenter, copy the **chain.pem** file, which contains all the certificates in the chain, to `/usr/local/avamarclient/bin` on the proxy.

3 Set proper operating system permissions on the certificate by typing:

chmod 600 /usr/local/avamarclient/bin/<*certificate_file*>

where, *certificate_file* is either a server certificate or a chain certificate.

4 Open `/usr/local/avamarclient/var/avvcbimageAll.cmd` in a UNIX text editor.

5 Append the following content at the end of the file:

--ssl_server_authentication_file=/usr/local/avamarclient/bin/<*certificate_file*>

where, *certificate_file* is either a server certificate or a chain certificate.

**NOTE** Use **chain.pem** for a chain vCenter SSL certificate.

6 Save the changes and close `avvcbimageAll.cmd`.

7 Open `/usr/local/avamarclient/var/avvmwfileAll.cmd` in a UNIX text editor.

8 Append the following content at the end of the file:

--ssl_server_authentication_file=/usr/local/avamarclient/bin/<*certificate_file*>

where, *certificate_file* is either a server certificate or a chain certificate.

**NOTE** Use **chain.pem** for a chain vCenter SSL certificate.

9 Save the changes and close `avvmwfileAll.cmd`.

10 Open `/etc/vmware/config` in a UNIX text editor.

11 Append the following contents at the end of the file:

vix.enableSslCertificateCheck = "true"

vix.sslCertificateFile = "/usr/local/avamarclient/bin/<*certificate_file*>"

where, *certificate_file* is either a server certificate or a chain certificate.

12 Save the changes and close `/etc/vmware/config`.

13 Open `/usr/local/avamarclient/var/vddkconfig.ini` in a UNIX text editor.

14 Change the value of **vixDiskLib.linuxSSL.verifyCertificates** from **0** to **1**.

15 Save the changes and close `vddkconfig.ini`.

16 Ensure that there are no running backup or restore jobs on the proxy.

17 Restart the `avagent` and `vmwareflr` services by running the following commands:

service avagent-vmware restart

service vmwareflr restart

### Result

The proxy can use and validate SSL certificates when connecting to the vCenter server.

Repeat this procedure for each proxy that you deploy in the vCenter environment.

# Configuring VDP 6

This chapter includes the following topics:


- "Login Security" on page 54
- "Configuring the Customer Experience Improvement Program" on page 54
- "Configuring and Monitoring" on page 55
- "Monitoring VDP Activity" on page 60
- "VDP Shutdown and Startup Procedures" on page 62

# Login Security

To increase security of the VDP configuration utility, after 5 unsuccessful login attempts from any user, the VDP configuration utility is locked for 5 minutes. This prevents any login attempts during that time. Any existing sessions that are already logged in are unaffected.

If a lockout occurs:

- The VDP configuration utility informs the user during a login attempt.
- A vCenter event about the lockout is generated.
- The `vdp-configure.log` captures the time, the user, the source address, and request headers of the failed login attempts that led up to the lockout.

  To view the `vdp-configure.log` file, log in to the VDP appliance by using terminal or console admin user and default password.

# Configuring the Customer Experience Improvement Program

The Customer Experience Improvement Program is an option that enables you to send encrypted configuration and usage information about the VDP environment to VMware servers for analysis.

**NOTE** The Customer Experience Improvement Program (PhoneHome) is not supported in a pure IPv6 vSphere Data Protection environment.

VDP sends the following types of data to VMware:

- Version information for VDP and vSphere
- VDP time zone and uptime
- Number of protected and unprotected VMs
- Amount of configured capacity
- Data Domain in use (true/false)
- Capacity utilization
- Number of internal and external proxies configured and registered
- Number of work orders per proxy configured by the user
- Maintenance and management status
- Latest valid checkpoint date
- Agents in use (SQL Server, Exchange, Exchange GLR/DAG, and SharePoint)
- Last integrity checkpoint status
- VDP state
- Details of job failures with timestamp and failure code
- Number of jobs that have failed against each failure code

During the installation of VDP, you can enable the Customer Experience Improvement Program from the Product Improvement page in the VDP Configuration Utility. The Product Improvement page provides the **Enable Customer Experience Improvement Program** checkbox. By default, the checkbox is not selected.

Anytime after the installation, you can either enable or disable the Customer Experience Improvement Program by accessing the VDP Configure Utility from a web browser.

### Prerequisites

In networks that are protected by a firewall, you may need to modify the network to help prevent connectivity issues when the Customer Experience Improvement Program attempts to upload data to the VMware server. To ensure that the firewall does not prevent Customer Experience Improvement Program from uploading data to the VMware server, open the network to the following VMware servers:

- https://vcsa.vmware.com:443
- https://phtransfer.vmware.com:443

### Procedure

1 Open a web browser and type:

**https://**<IP_address_VDP_appliance>**:8543/vdp-configure/**

2 Log in with the VDP username and password.

3 On the **Configuration** tab, click the icon and select **Product Improvement**.

The **Product Improvement** page appears.

4 Either enable or disable the Customer Experience Program by either selecting or clearing the **Enable Customer Experience Improvement Program** checkbox respectively.

# Configuring and Monitoring

From the vSphere web client, you can view and modify backup window configuration details, in addition to information about the appliance and storage. You can also configure the VDP appliance to send email reports on a scheduled basis.

## Viewing Backup Appliance Configuration

The backup appliance information provides information about backup appliance, storage summary, and backup windows configuration. The backup appliance information includes:

- Display name
- Product name
- IP Address
- Major Version (VDP version number)
- Minor Version (used by Technical Support)
- Status
- Host
- vCenter server
- VDP backup user
- VDP appliance time
- Time zone

These options are configured during the VDP appliance installation. They can be edited through the VDP Configure utility. "Post-Installation Configuration of VDP Appliance" on page 37 provides additional information.

The VDP appliance storage summary includes:

- **Capacity**: The total capacity of the VDP appliance.
- **Space free**: The amount of space that is available for backups.
- **Deduplicated size**: The amount of disk space that the backups are consuming in the deduplicated format.
- **Non-deduplicated size:** The amount of disk space that the backups would consume if they are converted to a native and non-deduplicated format.

**NOTE** To refresh the data on the page, click the refresh icon beside **Storage Summary**.

On the left, a color container variant is presented to check the VDP capacity:

- When the storage capacity used is under 80%, the container is green.
- When the storage capacity used is between 80% and 95%, the container is yellow. The following message appears:

  `The VDP storage is nearly full.`

- When the storage capacity used is between 95% and 100%, the container is red. The following message appears:

  `The VDP storage is nearly full.`

- When the storage capacity used is 100%, the container is red. The following message appears:

  `The VDP storage is full.`

**NOTE** If a Data Domain system is added to the VDP configuration, the Data Domain storage summary also appears.

Figure 6-4 illustrates the backup window configuration.

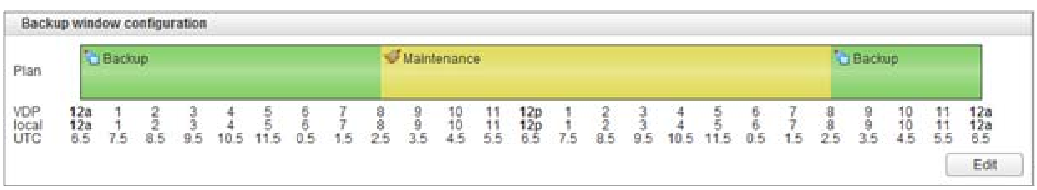



**Figure 6-4.** Backup window configuration

Each 24-hour day is divided into two operational windows:

- **Backup window:** The portion of each day reserved for performing normal scheduled backups.
- **Maintenance window:** The portion of each day reserved for performing routine VDP maintenance activities, such as integrity checks. Do not schedule backups or perform a "Backup Now" operation while VDP is in maintenance mode. The backup jobs will run but they will consume resources VDP needs for maintenance tasks.

  The jobs that are running when the maintenance window begins or that run during the maintenance window continue to run.

**NOTE** Because the blackout window is eliminated, activities such as, integrity checks and garbage collection now continuously run during the maintenance window.

## Editing the Backup Window

You can change the amount of time available for processing backup requests.

### Prerequisites

- VDP is installed and configured.
- You are logged in to the vSphere web client and connected to the VDP appliance.

### Procedure

1. Connect to the VDP appliance. "Accessing VDP" on page 106 provides information.
2. From the VDP user interface, select the **Configuration** tab (by default you are on the Backup appliance view).
3. At the bottom right corner of the Backup appliance view, click the **Edit** button.

   The backup start time and duration time options appear as shown in Figure 6-5:

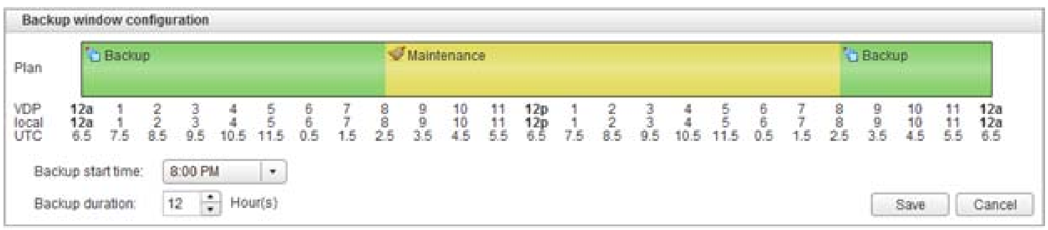



**Figure 6-5.** Backup window configuration in the edit mode

4. Use the drop down arrow to choose the start time for the backup window.
5. Enter the duration of the backup window. The minimum backup window is 4 hours and the maximum backup window is 16 hours.
6. Click **Save**.

   A dialog appears telling you that the settings were saved successfully.
7. Click **OK**.

VDP changes the backup window configuration.

## Configuring the Email Notifications and Reports

You can configure VDP to send SMTP email reports to specified recipients. If you have enabled email reports, the email messages that include the following information are sent:

- Items that need attention
- VDP appliance status
- Backup jobs summary
- Virtual Machines summary
- Replication summary
- Unprotected Virtual Machines

**NOTE** The VDP email notifications and reports do not support carbon copies (CCs), blind carbon copies (BCCs), and SSL certificates.

### Prerequisites

- VDP is installed and configured.
- You are logged in to the vSphere web client and connected to the VDP appliance.
- The email account for email notifications and reports must exist.

### Procedure

1. Connect to the VDP appliance. “Accessing VDP” on page 106 provides information.
2. From the VDP user interface, select the **Configuration** tab.
3. Select **Email** and click **Edit**.
4. Specify the following fields:
   - **Enable email reports**: Select this option to enable email reports.
   - **Outgoing mail server**: Type the name of the SMTP server that want to use to send email. You can type this name as an IP address, a host name, or a fully qualified domain name. The VDP appliance must be able to resolve the name that you have typed.

     The default port for non-authenticated email servers is 25. The default port of authenticated mail servers is 587. You can specify a different port by appending a colon (:) and the port number to the server name. For example, to specify the use of port 8025 on server “emailserver” enter emailserver:8025.
   - (Optional) **My server requires me to log in**: Select this option if your SMTP server requires authentication.
   - **Username**: Type the username that you want to authenticate with.
   - **Password**: Type the password that is associated with the username.

     VDP does not validate the password. The password that you type is passed directly to the email server.
   - **From address**: Type the email address from where you would like the email report. This can only be a single address.
   - **To address(es)**: Type a comma-separated list of up to 10 email addresses.
   - **Send time**: From the drop-down list, select the time you want VDP to email reports.
   - **Send day(s)**: Check the days you want the reports sent.
   - **Report Locale**: Select the country for the email reports.
   - **Enable CSV Attachment**: Select this option if you want to receive the email report as an attachment in the CSV format.
   - **Enable VDP alarm email notification**: Select this option to enable the VDP alarm email notification. If you select this option, VDP sends an email notification when the VDP alarm is triggered or if a posted event changes the VDP alarm status from yellow to red.
5. Click the **Save** button.
6. To test your email configuration, click **Send test email**.

## Viewing the User Interface Log

Clicking **Log** on the **Configuration** tab displays the user interface log for VDP. This is a high-level log that details the activities that have been initiated with the user interface and that identifies some key status items.

Click **Export View** to save the details that are displayed on the screen to file on the machine where your browser is running.

More detailed logs can be downloaded using the VDP configuration utility. Refer to "Collecting VDP Logs or Diagnostic Information" on page 40 for instructions.

## Running an Integrity Check

Integrity checks verify and maintain data integrity on the deduplication store. The output of an integrity check is a checkpoint. By default, VDP runs an integrity check every day during the maintenance window. In addition, you can start the integrity check manually.

**CAUTION** If the VDP appliance displays an alarm that the last valid integrity check failed or is out of date, run a manual integrity check. If you allow for the VDP appliance to continue to make backups while the integrity check is out of date, you are risking losing potential backup data if a rollback to the last validated checkpoint is ever required.

You can see a list of all of the VDP checkpoints through the VDP Configure utility, **Rollback** tab. See "Rolling Back an Appliance" on page 43 for additional information.

### Prerequisites

- VDP is installed and configured.
- You are logged in to the vSphere web client and connected to the VDP appliance.

### Procedure

1. Connect to the VDP appliance. "Accessing VDP" on page 106 provides information.
2. From the **Configuration** tab in the VDP user interface, click the  icon and select **Run integrity check**.

   A confirmation screen appears, asking if you want to perform a manual integrity check.
3. Click **Yes**.

   A message appears informing you that the integrity check has been initiated.
4. Click **OK**.

   VDP starts the integrity check.
5. Monitor the Integrity Check progress through Recent Tasks.

**NOTE** When the VDP Integrity Check is running, the Maintenance service is stopped. This may cause a temporary VDP error. Wait until the Integrity Check is complete and the Maintenance service automatically restarts and the VDP error message is resolved.

# Monitoring VDP Activity

You can monitor the activities of the VDP appliance by using one of the following methods:

- "Viewing Recent Tasks" on page 60.
- "Viewing Alarms" on page 61
- "Viewing the Event Console" on page 62

Tasks, events, and alarms that are generated by VDP are prefaced by "VDP:" Note, however, that some of the tasks and events that occur as part of VDP processes are performed by the vCenter server and do not have this prefix.

For example, if VDP runs a scheduled backup job against a running virtual machine, the following task entries are created:

1. Create virtual machine snapshot (vCenter acting on the virtual machine to be backed up)
2. VDP: Scheduled backup job (VDP starting the backup job)
3. Reconfigure virtual machine (the VDP appliance requesting services from virtual center)
4. Remove snapshot (virtual center acting on the virtual machine that has completed backing up)

To see only VDP-generated tasks or events in the Tasks or Events console, enter "VDP**:"** in the **Filter** field.

## Viewing Recent Tasks

VDP generates task entries in the Recent Tasks windows when it performs the following operations:

- Backups
- Automatic Backup Verification
- Restores
- Replications
- Replication Recovery
- VDP-Configuration
- Integrity Checks

Clicking on a task entry in the Recent Tasks window displays task details in the pane at the bottom of the screen. Task details can also be displayed by clicking the link next to the virtual machine icon in the **Running** tab under **Recent Tasks.**

You can cancel the tasks also from the **Running** tasks jobs panel by clicking the cancel task icon.

## Viewing Alarms

Table 6-10 lists the alarms that the vSphere Data Protection (VDP) appliance can trigger:

**Table 6-10.** vSphere Data Protection alarms

| Alarm Name | Alarm Description |
|---|---|
| VDP: [001] The most recent checkpoint for the VDP appliance is outdated. | From the **Configuration** tab of the VDP user interface, click the All Actions icon and select "Run integrity check." |
| VDP: [002] The VDP appliance is nearly full. | The VDP appliance is nearly out of space for additional backups. You can free the space on the appliance by manually deleting either unnecessary or older backups and by modifying the retention policies of the backup jobs to shorten the backup retention time. |
| VDP: [003] The VDP appliance is full. | The VDP storage is full. The appliance runs in the read-only mode till you provide additional space. You can free the space by deleting either unnecessary or older backups. |
| VDP: [004] The VDP appliance datastore is approaching maximum capacity. | The datastore where the VDP appliance provisioned its disks is approaching maximum capacity. When the maximum capacity of the datastore is reached, the VDP appliance will be suspended. The appliance cannot be resumed until additional space is made available on the datastore. |
| VDP: [005] Core services are not running. | Start Core services by using the VDP Configure utility. |
| VDP: [006] Management services are not running. | Start Management services by using the VDP Configure utility. |
| VDP: [007] File system services are not running (supported in VDP versions 5.5 and lower) | Start File system services by using the VDP Configure utility.<br>NOTE: This alarm is not supported in VDP version 5.8 or later. |
| VDP: [008] File level restore services are not running. | Start File level restore services by using the VDP Configure utility. |
| VDP: [009] Maintenance services are not running. | Start Maintenance services by using the VDP Configure utility. |
| VDP: [010] Backup scheduler is not running. | Start Backup scheduler by using the VDP Configure utility. |
| VDP: [013] Protected Virtual Machine limit exceeded. | The supported number of protected Virtual Machines has been exceeded. |
| VDP: [014] Backup Recovery services are not running. | Start Backup Recovery services by using the VDP Configure utility. |
| VDP: [015] Replication services are not running. | Start replication services by using the VDP Configure utility. |
| VDP:[016] The Data Domain storage is nearly full | The Data Domain storage is nearly full. You can free the space on the Data Domain storage by manually deleting the unnecessary or older backups, and modifying the retention policies of the backup jobs to shorten the backup retention time. |
| VDP:[017] The Data Domain storage is full | The Data Domain storage is full. You can free the space by manually deleting either unnecessary or older backups. |

## Viewing the Event Console

VDP can generate events of the following types: info, error, and warning. Examples of the following types of events are:

- **Info:** “VDP: Critical VMs Backup Job created.”
- **Warning:** “VDP: Unable to add Host123 client to backup job Critical VMs because. . .“
- **Error:** “VDP: Appliance has changed from Full Access to Read Only.”

VDP generates events on all state changes in the appliance. As a general rule, state changes that degrade the capabilities of the appliance are labeled errors, and state changes that improve the capabilities are labeled informational. For example, when starting an integrity check, VDP generates an event that is labeled an error because the appliance is set to read only before performing the integrity check. After the integrity check, VDP generates an event that is labeled informational because the appliance changes from read-only to full access.

Clicking on an event entry displays details of that event, which includes a link to **Show** related events.

## Persisted and Pop-up Messages to Show the Capacity Usage Issue

If either the VDP appliance storage or the Data Domain storage usage is at a critical point, an alert bar with a message appears on the vSphere web client. The text and the icon in the message depends on the amount of the storage usage. The same text and icon appears in the **Storage Summary** section on the **Configuration** tab.

You can either hide or show the alert bar by using the hide/show control, which exists below the alert bar. Hiding and showing of the alert bar is animated with fade and resize effects.

An icon that indicates a high severity message appears on the left side of the hide/show control. If you have hidden the alert bar, the high severity message icon appears to acknowledge you the capacity alerts.

Also, the alert bar contains the **Configuration - Backup Appliance** link on the right side. If you click the link, the **Backup Appliance** panel on the **Configuration** tab appears and the link on the alert bar is hidden.

If the storage usage is not critical, the alert bar and the hide/show control neither appear nor included in the layout.

If you resize the web browser window, the alert bar is scaled according to the content. A tool tip appears on the message if the text is truncated to fit in the bar.

# VDP Shutdown and Startup Procedures

If you need to shut down the VDP appliance, use the **Shut Down Guest OS** action from the vCenter server web client. This action automatically performs a clean shutdown of the appliance. If the appliance is powered off without the **Shut Down Guest OS** action, corruption might occur. It can take up to 30 minutes to shut down and restart the VDP appliance. You can monitor the status by using the virtual machine console. After an appliance is shut down, you can restart it by using the **Power On** action from the vCenter server web client.

If the appliance does not shut down properly, when it restarts it rolls back to the last validated checkpoint. This means any changes to backup jobs or backups that occur between the checkpoint and the unexpected shutdown are lost. This is expected behavior and is used to ensure system corruption does not occur from unexpected shutdowns. See “Rolling Back an Appliance” on page 43 for additional information.

The VDP appliance is designed to be run 24x7 to support maintenance operations and to be available for restore operations. Do not shut down the VDP appliance unless there is a specific reason for shutdown.

**NOTE** Before vCenter server patches or upgrades, use the VDP shutdown procedure.

# Proxies 7

This chapter includes the following topics:


- "Proxy Overview" on page 64
- "Managing Proxy Throughput" on page 65
- "External Proxy Support" on page 66
- "(Optional) Configuring Proxy Certificate Authentication" on page 68
- "Monitoring External Proxy Health Status" on page 69

# Proxy Overview

This chapter describes the use of internal and external proxies in VDP Appliances.

After initial deployment, a VDP appliance has only internal proxy configured. After the VDP appliance is fully deployed, you can deploy up to eight external proxies through the VDP Configuration UI. If you configure an external proxy for the VDP appliance, the internal proxy automatically is disabled during the process.

## Considerations Before Deploying an External Proxy

Consider the following when deploying one or more proxies for use with the VDP appliance:

- The VDP appliance does not have access to a datastore, preventing the HotAdd transport method from being used. Refer to the following VMware knowledgebase article to define the VMware HotAdd best practices from a backup point of view:

  http://kb.vmware.com/kb/1008072

- The need to perform more backups concurrently while not constrained by resources on the vSphere host nor by datastore performance. A maximum of 24 concurrent backup jobs are supported.
- To perform File Level Recovery (FLR) on Logical Volume Manager (LVM) or EXT4 based file systems, requires an external proxy.

## Deployment of External Proxies

To take advantage of HotAdd during the backup process, the proxy appliance must have direct access to the datastore where the target VM is located. The backup agent appliance leverages the datastore access from the vSphere host to which it is associated. When deploying the external proxy, verify that the appliance has access to the desired datastores by way of the vSphere host.

If the proxy appliance does not have access to the datastore where the target VM is located, the NBDtransport method will be invoked instead of HotAdd, which may greatly reduce the speed of the backup.

## Number of Proxies to Deploy and Proxy Throughputs per Proxy

When considering the number of proxies to deploy and the configuration of the number of proxy throughputs per proxy, consider the following as best practices:

### When using a single proxy

- For an external proxy, if deployed with the default configuration for memory and CPUs, the optimal number of proxy throughputs is six when running backups. Increasing the proxy throughputs per proxy beyond six leads to performance degradation.
- Increasing the CPU on the external proxy results in better performance than increasing the memory on the single external proxy.
- When the external proxy configuration is modified to eight CPUs, running eight proxy throughputs per proxy is optimal. This is true for Level 0 and Level 1 backups. With this configuration, network bandwidth becomes the limiting factor.

### When using multiple external proxies for Level 0 backups

- Maximizing the number of proxies deployed may not produce the best performance results.
- As you increase the number of proxies, the optimal number of proxy throughputs per proxy may decrease. With two external proxies running, for example, you may see the best results while running six proxy throughputs per proxy. With four external proxies running, you may see the best results while running four proxy throughputs per proxy. This may be constrained by the number of proxies per vSphere host.

- It is better to increase the number of proxy throughputs per proxy instead of the number of proxies when performing a set number of backups.

### When using multiple external proxies for Level 1 backups

When running incremental (Level 1) backups for a virtual machine, consider increasing the number of proxies. Running four proxy throughputs per proxy provides better performance than fewer proxies running eight proxy throughputs per proxy.

### Increasing the volume of backups that run on the VDP appliance

After backups have been running, you may need to increase the volume of the backups performed during a specific timeframe by the VDP appliance. When increasing the quantity of backups performed in a given timeframe, consider the following best practices:

- If the vSphere server resources are a constraint, run more proxy throughputs per proxy and reduce the number of proxies.
- If the vSphere server resources are not constrained, increase the number of proxies and maintain four proxy throughputs per proxy.
- If six to eight external proxies are necessary to process the desired backups, increase the number of proxy throughputs per proxy and limit the number of external proxies deployed.

### Reducing concurrent backups run on the VDP appliance

After backups have been running on the VDP appliance, you may need to reduce the number of concurrent backups being run by the VDP appliance to limit the load on the datastores and associated storage. When reducing the quantity of concurrent backups being run by the VDP appliance, consider the following best practices:

- Place the VDP appliance on a datastore with faster write performance, preferring iSCSI to NFS for the datastore.
- If there is high load on the datastore where the protected virtual machines reside during backups, reduce the number of proxies and the number of proxy throughputs to proxy to four or less. This will reduce the quantity of seek and read operations when performing backups.

## Managing Proxy Throughput

The Manage Proxy Throughput wizard allows you to configure the number of simultaneous backups and restores, based on your infrastructure. You can use the Manage Proxy Throughput wizard to set values for both internal and external proxies. The VDP appliance supports up to 8 external proxies, and the maximum number of concurrent backup jobs is 24.

For more information, refer to "Number of Proxies to Deploy and Proxy Throughputs per Proxy" on page 64.

CAUTION The number you select to simultaneously back up and restore clients is a global setting. This setting applies to all internal and external proxy settings.

### Procedure

1. Open a web browser and type:

   **https://**<*IP_address_VDP_Appliance*>**:8543/vdp-configure/**

2. Log in with the VDP username and password.
3. Click the **Configuration** tab.
4. From the **Action** list, select **Manage Proxy Throughput**.

   The Manage Proxy page appears.

5. Select the number, from 1 to 8, of backup and restore clients that you want to run simultaneously.

6 Click **Next**.

The Ready to Complete page appears.

7 Click **Finish** to apply the changes.

# External Proxy Support

VDP Advanced appliance versions 5.5.6 and earlier are configured with internal proxies only, where the proxy services run inside the VDP appliance and are registered to manage job requests from the appliance. External proxies can be configured for VDP 6.0 Appliances.

**NOTE** The VDP appliance with external proxies supports File Level Recovery (FLR) on virtual machines with EXT4 file systems.

## Best Practice for External Proxy Deployment

**NOTE** If configuring the appliance to use external proxies, the internal proxies are automatically disabled.

Using HotAdd transport is usually faster than over-the-network backup and restore using NBD (network block device) transport.

## External Proxy Best Practices

- **Clean up orphan proxies**

  External proxies that are registered to the VDP appliance but no longer exist in the vCenter inventory are considered orphan proxies. If the proxy virtual machine remains in the vCenter inventory and you see the "Either the VM is deleted or is currently not managed by vCenter" warning for the proxy, you can restart the proxy virtual machine through the VDP configuration utility. If the problem persists, the issue might be the hostname is not resolvable due to incorrect network configuration. To work around this issue, delete the orphan proxy and re-deploy a new proxy.

- Remove the proxy's ESXi host from the vCenter

  When an external proxy's ESXi host is removed from the vCenter, the VDP appliance considers the proxy as an orphan and displays the "Either the VM is deleted or is currently managed by vCenter" warning when it is selected in the VDP configuration utility. Unless the ESXi host will be added back to the vCenter, delete the orphan proxy entry from the VDP configuration utility.

- **vCenter switch**

  If a proxy's ESXi host is not moved to the new vCenter along with the VDP's ESXi host, that proxy appears as an orphan in the VDP configuration utility. Unless you are planning to move the proxy's host to the new vCenter or move the VDP's host back to the original vCenter, delete the orphan proxy from VDP configuration utility.

- **VDP rollback**

  After the VDP appliance rolls back to a previous checkpoint, any external proxies added or deleted after the checkpoint was taken appear as orphan proxies in the VDP configuration utility. For added external proxies, you can restart them through the VDP configuration utility to re-register or redeploy them. For the deleted proxies, delete the orphan proxy entries from the VDP configuration utility.

- **Rebuild VDP from scratch and rollback**

  If the VDP appliance experienced issues that required a rebuild from scratch, and the appliance rolled back to a checkpoint from Data Domain, use the VDP configuration utility to change the appliance's password immediately. Any external proxies that were deployed before the checkpoint was taken are displayed as orphans. Delete these orphans from the VDP configuration utility. Any external proxies that were deployed on the rebuilt VDP appliance before the checkpoint rollback appear with a "Failed to authenticate with the proxy" warning. Update the proxy's password after you change the VDP appliance's password.

- **VDP disaster recovery from tape**

  If the VDP appliance is recovered from tape after a disaster, delete the orphan external proxies from the VDP configuration utility and deploy new ones.

- **Update password**

  When the VDP appliance's password is changed through the VDP configuration utility, the VDP appliance updates the password on all registered external proxies. If a proxy's password is not updated, you will see a "Failed to authenticate with the proxy" warning when the proxy is selected. You can update the proxy's password manually through the VDP configuration utility.

- **Restart proxy**

  If any previously-registered external proxies appear as not registered in the VDP configuration utility, restart the proxies, which will power cycle the proxies and reregister them to the VDP appliance.

- **Emergency restore**

  When an emergency restore operation is performed with external proxies, the internal proxy is automatically enabled. When the emergency restore operation completes and all the ESXi hosts are reconnected to the vCenter, you can either delete the internal proxy or delete all the external proxies. Do not enable both the internal and external proxies for non-emergency restore activities.

### Limitations

- The external VDP proxies are supported only on vSphere host 5.1 and later.
- The limit of external proxies per VDP appliance is 8.
- The maximum number of concurrent backup jobs is 24.
- File Level Recovery (FLR) in virtual machines that use the EXT4 file system requires the use of an external proxy.

## Adding an External Proxy

The Add Proxy wizard allows you to add and deploy up to 8 external proxies and register them with the VDP appliance from the VDP configuration utility.

### Procedure

1. Open a web browser and type:

   **https://**<*IP_address_VDP_Appliance*>**:8543/vdp-configure/**

2. Log in with the VDP username and password.
3. Click the **Configuration** tab.
4. From the **Action** list, select **Add Proxy**.

   The Add Proxy page appears.

5. On the Host and Storage page, specify the following credentials, and then click **Next**.

   - **Host**: Select a destination host from the list.
   - **Storage**: Select a destination storage device from the list.
   - **Network Connection**: Select a network connection for the virtual machine from the list.

6 On the Network page, specify the following network settings, and then click **Next**.

- **IPv4 static address**: The Internet Protocol v4 or v6 setting of the interface.
- **Netmask**: The network mask of the IPv4 or IPv6 static address.
- **Gateway**: The gateway address of the IPv4 or IPv6 static address.
- **Primary DNS**: The primary domain name system used for DNS resolution.
- **Secondary DNS**: The secondary domain name system used for DNS resolution.

7 On the Ready to Complete page, review the deployment settings.

**NOTE** You can change the name of the external proxy virtual machine on the Ready to Complete page.

8 Click **Finish** to deploy the proxy.

Ensure that the proxy deploys to the datastore that was selected in Step 5. After the successful deployment, if the proxy deploys to a VMware vSphere Distributed Resource Scheduler™ (DRS) cluster, the cluster can move the proxy by using storage vMotion. The jobs that are running on the proxy while it is migrating to a different storage are at risk. HotAdd does not work for the proxies that are located in a DRS cluster. Therefore, you must manually disable DRS for the proxy. This scenario applies to the VDP appliance also, if an internal proxy is in use.

Perform the following steps to manually disable DRS for the proxy:

1 Select the **DRS** tab on the cluster that hosts the proxy.

2 Click **Virtual Machine Options**.

3 Under the **Automation** level, locate the proxy virtual machine and change it to **Disable**.

## Disabling the Internal Proxy

When a user performs a rollback operation as a checkpoint performed during an external proxy backup, after the rollback completes, the VDP configuration utility displays the internal proxy, along with a warning message prompting the user to disable the internal proxy.

### Procedure

1 In the VDP configuration utility, select the internal proxy from the **Proxies** list.

2 Select **Manage proxy** from the Proxies Actions icon.

The Manage internal proxy wizard displays.

3 Enable the **Disable internal proxy** checkbox, and click **Next**.

4 On the Ready to Complete page of the wizard, click **Finish** to apply the changes.

# (Optional) Configuring Proxy Certificate Authentication

By default, proxies do not validate SSL certificates when connecting to the vCenter server. This can leave the vCenter server vulnerable to a man-in-the-middle exploitation, which might result in unauthorized access to the vCenter server. Configure each proxy to use an SSL certificate authentication when connecting to the vCenter server corrects this vulnerability. Chapter 5, "Securing Communication between VDP and vCenter," on page 47 provides more information.

# Monitoring External Proxy Health Status

## Health Status Criteria

The health status reported for the external proxy is based on the following criteria:

- Disk Usage
  - **Warning**: Any file system that is greater than 70% capacity.
  - **Critical Warning**: Any file system that is greater than 90% capacity.
- CPU Load
  - **Warning**: 15 minute load average is greater than 1.5.
  - **Critical Warning**: 15 minute load average is greater than 5.0.
- Memory Usage
  - **Warning**: Usage is greater than 85%.
  - **Critical Warning**: Usage is greater than 95%.

## External Proxy Logs

External proxy logs are visible only if at least one external proxy is configured. Downloading this log bundle streams a `.zip` file to the browser. The `.zip` file contains all the external proxy logs. The internal proxies are excluded from this log bundle.

For information about downloading logs, refer to "Collecting VDP Logs or Diagnostic Information" on page 40.

# Storage Management 8

This chapter contains the following topics:


- "Creating New Storage" on page 72
- "Attaching Existing VDP Disks" on page 74
- "Detaching and Reattaching Storage" on page 75
- "Viewing the Storage Configuration" on page 76

## Creating New Storage

The Initial Configuration wizard guides you through the storage type selection, device allocation on VDP storage disks, and the option to run the performance assessment tool.

### Limitations

You cannot migrate the appliance to a new host or to a new datastore while you are creating new storage.

### Prerequisites

- The VDP appliance is deployed and you are logged into the **Create Storage** page of the **Initial Configuration** wizard.
- After the OVA deployment, immediately disable vSphere HA on the VDP appliance. After performing the initial configuration on the VDP appliance, which includes creating a storage, you can enable vSphere HA for the appliance.

### Procedure

1. On the **Create Storage** page of the **Initial Configuration** wizard, select **Create new storage**. When you create a storage, the process creates a new storage node on the selected datastores.
2. Select one of the following capacity options and click **Next**.
   - 0.5
   - 1
   - 2
   - 4
   - 6
   - 8

   The Device Allocation page appears. When you create new storage, the number of required disks is known.
3. Select the provision type from the **Provision** drop-down list.
   - **Thick Lazy-Zeroed** (the default and recommended provision type): Thick lazy zeroed provisioning creates a virtual disk in a default thick format. Space required for the virtual disk is allocated when the virtual disk is created. Data remaining on the physical device is not erased during creation, but is zeroed out on demand at a later time on first write from the virtual machine.
   - **Thick Eager-Zeroed**: Thick eager zeroed provisioning creates a type of thick virtual disk that is used when data security is a concern. Space required for the virtual disk is allocated when the virtual disk is created. When you create a virtual disk using thick eager zeroed provisioning on a datastore that had previous data, the previous data is erased and cannot be recovered. It might take much longer to create disks in this format than to create other types of disks.
   - **Thin**: For the thin disk, you provision as much datastore space as the disk would require based on the value that you enter for the disk size. The thin disk starts small and uses only as much datastore space as the disk needs for its initial operations.
4. After all of the disks have been allocated to datastores, click **Next**.

   On the **Ready to Complete** page, you can run a performance analysis on the storage configuration and click **Next** to apply the changes. Though you can bypass the performance analysis test, you are strongly encouraged to run it.

5 When you click **Next**, you are prompted with a warning that the storage configuration will start and cannot be undone. Click **Yes** to continue.

Possible results are Passed, Failed, and Conditionally Passed if the seek test is selected (Excluded by Default). If all tests succeed, the result is Passed. If the write or read tests are unsuccessful, the result is Failed. If the seek test is selected and if the write and read tests are successful but the seek test fails, the result is Conditionally Passed.

a To run the test, select **Run performance analysis on storage configuration** to ensure that the storage configuration meets minimum performance expectations. Table 8-11 lists the minimum performance expectations.

This test performs write, read, and seek performance tests on the disks. Depending on your storage configuration, performance analysis can take from 30 minutes to several hours to complete.

b Select **Restart the appliance if successful** to automatically restart the appliance after the test succeeds and click **Next** to start the test.

The performance analysis test is server-initiated and you can close the browser while the test runs.

- If the test succeeds, a message that states that the configuration is complete appears and the server automatically restarts the appliance.
- If the test conditionally passes or fails, the results of the performance analysis appear, but the server does not automatically restart the appliance. To view the test results, you must relog in to VDP-Configure and manually restart the client.

**NOTE** If you do not click **Restart** within 59 seconds, the appliance automatically restarts and starts the services. After the VDP appliance restarts, it performs a series of automated configuration steps. These steps can take 30-45 minutes or more to complete.

The appliance automatically restarts under the following circumstances:

- If you opt to perform performance analysis at the end of the VDP Configuration wizard.
- In the VDP Configuration wizard, you selected the checkbox to restart the VDP appliance after the tests succeed.
- After the tests succeed.

## Minimum Storage Performance

When configuring the VDP appliance, the performance test produces varying results depending on the size of the appliance being deployed.

Table 8-11 lists the minimum expectations for Read, Write, and Seek performance by VDP appliance size.

**Table 8-11.** Minimum expectations for storage performance

| VDP Appliance Size (in TB) | Disk Size | Read Minimum | Write Minimum | Seek Minimum |
|---|---|---|---|---|
| Single datastore | 256 GB | 10 MB/s | 20 MB/s | 106 Seeks/s |
| 0.5 | 256 GB | 60 MB/s | 30 MB/s | 400 Seeks/s |
| 1.0 | 512 GB | 60 MB/s | 30 MB/s | 400 Seeks/s |
| 2.0 | 1024 GB | 60 MB/s | 30 MB/s | 400 Seeks/s |
| 4.0 | 1024 GB | 80 MB/s | 40 MB/s | 400 Seeks/s |
| 6.0 | 1024 GB | 80 MB/s | 40 MB/s | 400 Seeks/s |
| 8.0 | 1024 GB | 150 MB/s | 120 MB/s | 400 Seeks/s |

# Attaching Existing VDP Disks

The Create Storage page in the Initial Configuration wizard provides the **Attach existing VDP disks** option. This option enables you to browse the datastores and select the previously-used VDP disks, and then proceeds to automatically attach the selected disks to the new VDP appliance. These disks are automatically added to the new VDP appliance.

CAUTION Do not attempt to manually attach the used VDP disks to the new VDP appliance without following the steps in this procedure. If you import the used VDP disks, only the restore points are imported, but not the backup jobs and the configuration information such as, email reports. You must recreate the jobs. The imported restore points exist in **/REPLICATE/VDP_IMPORTS** domain. On the vSphere web client, you can see the actual restore points in **/REPLICATE/VDP_IMPORTS** domain. If VDP detects an incomplete or invalid storage configuration, the operation fails.

When you attach existing storage, you do not need to select a capacity option, as you are required to do when creating new storage.

The following changes occur when previously-used VDP disks are attached to the new VDP appliance:

- All backup jobs associated with the previously-used VDP are deleted and must be re-created.
- All replication jobs associated with the previously-used VDP are deleted and must be re-created.
- Restore points that are associated with the previously-used VDP are kept intact. The restore points appear with the original VM name appended with a string of random letters.
- On the Set Restore Options page of the Restore a backup wizard, the **Restore to Original Location** option is disabled for the restore points associated with the previously-used VDP disks.
- Email reporting must be reconfigured.

### Prerequisites

- Before you can attach existing storage, you must install and configure the VDP appliance described in "VDP Installation and Configuration" on page 21.
- Ensure that you back up all the VDP storage which you intend on attaching to the VDP appliance.

### Procedure

1 On the Create Storage page of the Initial Configuration wizard, select **Attach existing VDP disks** and click **Next**.

   The Device Allocation dialog box appears.

2 Click the first ellipsis button:

   a Browse to the first vmdk file you want to attach.

   b Highlight the vmdk file and click **Select**.

NOTE You can select to attach only the data disks of the previously-used VDP appliance. You cannot select the OS boot partition. If you select the primary disk partition of 200 GB (the OS boot partition), an error message appears.

   You can attach vmdk files in any order. After you have selected the first vmdk file, the system analyzes the disk and defines how many additional disks should be selected.

NOTE At any time during the attach process, you can click **Reset** to reset the Device Allocation dialog box to its original state.

3 Click the ellipsis button that corresponds to the next disk to be defined:

  a Browse to the next vmdk file you want to attach.

  b Highlight the vmdk file and click **Select**.

  Each disk is validated as a matching disk before it is added. If the validation fails, an error message appears. Hover over the red highlighted disk to see the error message.

4 Repeat Step 3 for all remaining disks.

5 After all disks have been allocated, click **Next** to validate the complete set of disks.

  The Ready to Complete page appears.

6 Click **Next**.

  The system displays the message:

  The following process will start the storage configuration. This cannot be undone. Do you wish to continue?

7 Click **Yes**.

  The system prompts you to provide the root password associated with the previously-used VDP appliance.

8 Type the root password of the previously-used VDP appliance into the **Password** and **Verify Password** text boxes, and then click **OK**.

9 Click **Finish** to apply the changes and reboot.

**NOTE** After a successful storage configuration, the system automatically reboots and starts the services. After the VDP appliance reboots, it performs a series of automated configuration steps. These steps can take 30-45 minutes or more to complete.

After the configuration is complete, an integrity check begins.

# Detaching and Reattaching Storage

The following procedure explains the steps you must perform if the primary disk partition (the OS boot partition) on the VDP appliance becomes corrupt or lost, resulting in an unrecoverable VDP appliance.

## Prerequisites

- At least one validated checkpoint is present on the VDP appliance where the vmdk files are being detached and reattached.
- A new VDP appliance is deployed which is compatible with the older VMDK disk data (the VDP appliance must be the same version as the disk data or a newer version).
- The vmdk files from the previous VDP appliance must be on a datastore that is accessible by the newly-deployed VDP appliance.

**NOTE** During the reattach process, you will be prompted to enter the root password for the old VDP appliance.

## Best Practices

- Make a backup copy of all vmdk files before reattaching them to a VDP appliance.
- If possible, detach the vmdk files from the VDP appliance after shutting down the VDP appliance by using the **Shut Down Guest OS** action. Otherwise, as a last resort, power off the VM.
- Before detaching a vmdk file from the VDP appliance, note the full path and name of the vmdk file. You will need this information when reattaching the disk to the newly-deployed VDP appliance.

### Procedure

1 In the vSphere Client, navigate to the VDP appliance and perform a **Shut Down Guest OS** action on the virtual machine.

NOTE If the **Shut Down Guest OS** action is grayed out, navigate to **vCenter > Hosts and Clusters**, right-click the VDP appliance, and select **Power off VM**.

2 Detach the vmdk files from the VDP appliance:

   a From the vSphere Web client, log in as a user who has privileges to edit hardware settings.

   b Navigate to **vCenter > Hosts and Clusters.**

   c In the tree on the left, click the disclosure arrows until the VDP appliance appears.

   d Right-click the VDP appliance and select **Edit Settings**.

   The virtual machine properties appear. The **Hardware** tab is selected by default.

   Hard disk 1 is always the primary, 200 GB OS boot partition. Do not remove Hard disk 1 from the VDP appliance.

   e Click Hard disk 2 from the list.

   f From the **Disk File** field, note the full path and name of the vmdk file. You will need this information when reattaching the disk.

   g Click **Remove**.

   h Under Removal Options, select **Remove from virtual machine**.

   i Repeat the removal option for each hard disk (2 through *x*) in the list.

   j After removing hard disks 2 through *x*, click **OK**.

3 On the Create Storage page of the Initial Configuration wizard, select **Attach existing VDP disks**, and follow the steps in "Attaching Existing VDP Disks" on page 74.

## Viewing the Storage Configuration

The **Storage** tab provides a storage summary, capacity utilization information, and performance analysis details.

The **Capacity Utilization** button displays a page that provides status information for the datastore:

- A horizontal bar graph next to the datastore icon shows how full the datastore is as a percentage.
- A pie chart shows a breakdown of the space in the datastore. Orange represents the amount of space used by the datastore. Green represents the amount free space on the datastore. Blue represents the space on the datastore that is used by other applications that run on virtual machines deployed on the datastore.
- A table next to the pie chart lists each data partition name, size, provision type, and vmdk file name. The following example shows information for a data partition named data 01:

| data 01 | 256GiB | Thin | sample-vdp-241168_6..0.0.117_1.vmdk |
|---|---|---|---|

  In the example, 256 GiB indicates the maximum size that can be utilized.

The **Performance Analysis** button displays a table that contains statistics from a performance analysis test.

| Column | Description |
|---|---|
| Datastore | Name of the datastore. |
| Completed On | Date that the test completed. |
| Result | A test can display one of the following results:<br>■ Never Run<br>■ Passed<br>■ Failed<br>If the write or read tests are unsuccessful, the result is “Failed.”<br>■ Conditionally Passed<br>If the write and read tests are successful, but the seek test (Optional) failed, the result is “Conditionally passed.” |
| Read (MiB/s) | Megabytes per second for the read test. |
| Write (MiB/s) | Megabytes per second for the write test. |
| Seek (Seeks/s) | Megabytes per second for the seek test. |

### Prerequisites

The VDP storage disks are distributed across the available datastore locations, the disks are validated, the system is rebooted, and the system services are up and running.

### Procedure

1 Log in to the VDP-Configure URL:

**https://**<*IP_address_VDP_appliance*>**:8543/vdp-configure/**

2 Click the **Storage** tab.

The Storage Summary appears, displaying the available datastores and the amount of total usable storage and storage capacity that is available for each datastore.

**NOTE** To refresh the data on the page, click the refresh icon beside **Storage Summary**.

3 To view status information about the datastore, click **Capacity Utilization**. This page is the default.

4 To run a performance test, click **Performance Analysis**, select a datastore in the table, and then click **Run**.

The performance analysis test creates a 256 GB vmdk on the datastore, and runs read, write, and seek tests to check the datastore performance.

# Enabling the Seek Test

The Seek test was designed to get Seek performance of desired storage. This is mainly useful where underlying storage is not thin-provisioned. As most of the field implementations are using thin-provisioned volumes, Seek test results can be confusing. The Seek test is not suitable for virtual environments; hence, this test is excluded from default tests that are run as part of the DAT Test Suite. The code will still be present and the Seek test can be performed if the administrator includes it in the configuration file.

If you specifically want to run the Seek test, you must first change the configuration file using the instructions below. Then you can run the Performance Analysis Tool which includes the Seek test.

## Changing the Configuration File to Enable the Seek Test

1 Using PuTTY, SSH to the VDP appliance.

2 Open the `/usr/local/vdr/etc/benchmark-settings.xml` file, which requires write permissions.

3 Under Options, change <runSeekTest>false</runSeekTest> to <runSeekTest>true</runSeekTest>.

4 Save the file and exit PuTTY.

5 Run the Performance Analysis suite.

# Data Domain Integration 9

This chapter contains the following topics:


- "VDP and Data Domain System Integration" on page 80
- "Architecture Overview" on page 80
- "VDP Client Support" on page 80
- "Best Practices" on page 81
- "Pre-integration Requirements" on page 82
- "Preparing the Data Domain system for VDP Integration" on page 84
- "Adding a Data Domain System" on page 85
- "Editing the Data Domain System" on page 87
- "Deleting the Data Domain System from the VDP Appliance" on page 87
- "Backups with VDP and Data Domain" on page 89
- "Replication Control" on page 90
- "Server Maintenance Activity Monitoring" on page 91
- "Restoring Avamar Checkpoint backups from Data Domain systems" on page 92
- "Monitoring Data Domain from the VDP Appliance" on page 94
- "Reclaiming Storage on a Full Data Domain System" on page 95
- "Common Problems and Solutions" on page 97

# VDP and Data Domain System Integration

VDP and Data Domain system integration enables the following:

- The use of physical Data Domain systems as a backup target for VDP backups
- The target destination of backup data, which is set during the creation of a backup job
- Transparent user interaction to the backup destination (VDP or Data Domain)

# Architecture Overview

A Data Domain system performs deduplication through DD OS software. VDP source based deduplication to a Data Domain system is facilitated through the use of the DD Boost library.

VDP uses the DD Boost library through API-based integration to access and manipulate directories, files, and so forth, contained on the Data Domain file system. The DD Boost API gives VDP visibility into some of the properties and capabilities of the Data Domain system. This enables VDP to control backup images stored on Data Domain systems. It also enables VDP to manage maintenance activities and to control replication to remote Data Domain systems.

DD Boost is automatically installed on the VDP appliance during the addition of a Data Domain system.

Figure 9-6 depicts a high-level architecture of the combined VDP and Data Domain solution. With VDP and Data Domain integration, you can specify whether a specific backup policy targets a VDP appliance or a Data Domain system.

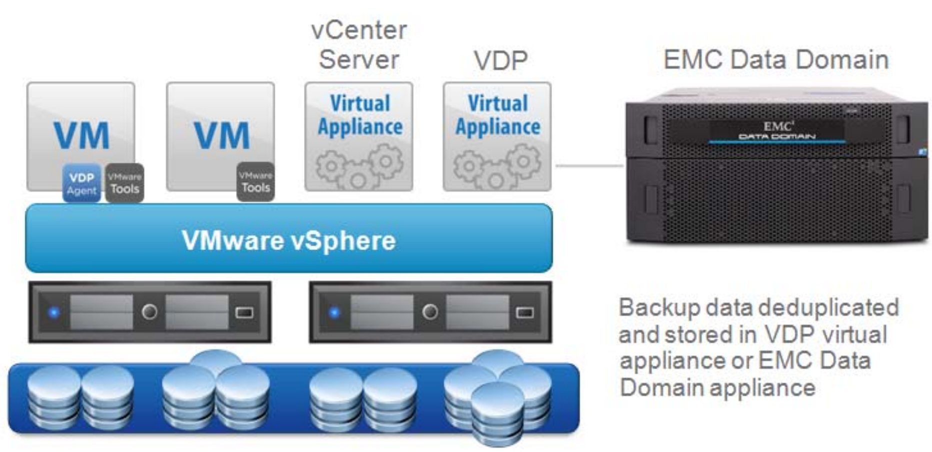



**Figure 9-6.** VDP and Data Domain solution

When you select the VDP appliance as the target for backup storage, the VDP appliance performs deduplication segment processing.

When you select a Data Domain system as the backup target, backup data is transferred to the Data Domain system. The related metadata that is generated is simultaneously sent to the VDP appliance for storage. The metadata enables the VDP appliance to perform restore operations from the Data Domain system.

# VDP Client Support

VDP and Data Domain system integration supports the following plug-ins:

- VDP Plug-in for Exchange Server VSS
- VDP Plug-in for SharePoint Server VSS
- VDP Plug-in for SQL Server

# Best Practices

### What are the VDP limitations with a Data Domain system attached?

VMware suggests protecting up to 25 virtual machines per TB of capacity on a VDP appliance. This variable is dependent upon the size of the virtual machines, the typical change rate, and the amount of data on each virtual machine. With these considerations, you can protect up to 200 virtual machines per VDP appliance backing up to a Data Domain system.

Because the backup data is stored on the Data Domain system, and only the backup job metadata is stored on the VDP appliance, deploy a 0.5 TB VDP appliance for a standard Data Domain system and a 1 TB VDP appliance for a 64 TB Data Domain system.

The following list provides the suggested number of VDP appliances deployed per Data Domain system:

- 1 VDP per DD160 and DD620
- 1 VDP per DD2200
- 2 VDP per DD2500 and DD4xxx
- 3 VDP per DD7200 and DD990

### What size VDP appliance do I need if I want all my backups to go to Data Domain?

With a Data Domain system attached to a VDP appliance as the storage device, the VDP appliance only stores the metadata for backups using the Data Domain system as the destination. It has been determined that a 16 TB Data Domain system only requires a 0.5 TB VDP appliance if all of the backup data is sent to the Data Domain system. If backups are going to be sent to the VDP appliance as well, the size of the VDP appliance should be increased accordingly based on the data to be stored on the VDP appliance. With a 64 TB or larger Data Domain system, you can deploy a 1 TB VDP appliance per 64 TB of Data Domain system storage you expect to consume with the backup data.

### I have many images, pictures, and PDF files in my VMs. Should I set the VDP appliance or the Data Domain system as the destination for these backups?

The Data Domain system provides better deduplication on images, pictures, and PDF files than the standard VDP appliance deduplication algorithms.

## Data Domain Limitations

The following are current limitations defined for the use of a Data Domain system with a VDP appliance:

- If a VDP appliance has a Data Domain system attached to it, the "Import Existing storage" functionality of the VDP appliance cannot be used for the VMDKs of the VDP appliance with the Data Domain system attached.
- Only one Data Domain system can be attached to a VDP appliance at a time.
- The Data Domain system cannot be deleted from the VDP Configuration UI. Use the manual steps defined in "Deleting the Data Domain System from the VDP Appliance" on page 87 to delete a Data Domain system.
- If the Data Domain and VDP connection is broken, the VDP appliance does not monitor the Data Domain system. Behaviors that might indicate the connection between the appliances are broken include, but are not limited to, the integrity check, `hfscheck` or backups failing.
- The Data Domain system or VDP appliance cannot be upgraded if the connection between them is broken.

## Backup

During a backup, the VDP appliance generates a backup request for the backup destination. If the backup request includes the option to use a Data Domain system as the destination, backup data is stored on the Data Domain system. Metadata is stored on the VDP appliance.

### Restore

The process of data recovery is transparent to the backup administrator. The backup administrator uses the same VDP recovery processes that are native to current VDP implementations.

### Security - Encryption

When using a VDP appliance with a Data Domain system attached, there are two potential backup data streams. If the backup data is being written to the VDP appliance, the backup data is always compressed and encrypted during flight. If the backup data is being routed to the Data Domain system, the `ddboost` utility encrypts the backup data as it is transmitted over the network to the Data Domain system.

### Data Migration

You cannot migrate backup data directly from the VDP appliance to the Data Domain system.

To use the Data Domain system as the backup target for backing up a VM or appliance instead of the VDP appliance, edit the backup job and define the destination as the Data Domain system, and then start performing backups to the Data Domain system. When you change the backup target to the Data Domain system, the next backup runs as full backup.

After you successfully perform a backup to the Data Domain system, you can delete the earlier backups from the VDP appliance. Refer to "Deleting a Backup Job" on page 119 for information on how to manually delete backups. If you do not manually delete backups, they expire without intervention.

## Pre-integration Requirements

Before you integrate a Data Domain system with a VDP appliance, review the following topics:


- "Network Throughput" on page 82
- "Network Configuration" on page 83
- "NTP Configuration" on page 83
- "Licensing" on page 83
- "Port Usage and Firewall Requirements" on page 83
- "Capacity" on page 83
- "Data Domain System Streams" on page 84
- "Existing Backup Products in Use with Data Domain" on page 84


**NOTE** This section assumes the VDP appliance and the Data Domain system are installed and configured.

### Network Throughput

With VDP, the VDP appliance and Data Domain systems can connect over a Wide Area Network (WAN). Before using this configuration, validate the firewall port requirements of the Data Domain system. "VDP Port Usage" on page 191 provides information.

You can use VDP appliance replication over a WAN to replicate data from source VDP appliance and Data Domain systems to target VDP appliances, provided they also have a Data Domain system attached.

Before integrating a Data Domain system with an VDP appliance, ensure that enough network bandwidth is available. To obtain the maximum throughput available on a Data Domain system (for restores, level zero backups, and subsequent incremental backups after a level-zero backup), verify that the network infrastructure provides more bandwidth than the bandwidth required by the maximum throughput of the Data Domain system. To see the network throughput, use the `system show performance` command on your Data Domain system:

```
system show performance [{hr | min | sec} [{hr | min | sec}]]
```

For example:

```
system show performance 24 hr 10 min
```

This command shows the system performance for the last 24 hours at 10 minute intervals. 1 minute is the minimum interval.

## Network Configuration

Configure (or verify) the following network configuration:

- Assign a Fully Qualified Domain Name (FQDN) to the Data Domain system.
- Use a FQDN (not an IP address) when registering a Data Domain system. This can limit the ability to route optimized duplication traffic exclusively through a registered interface.
- Ensure that DNS on the Data Domain system is properly configured.
- Ensure forward and reverse DNS lookups work between the following systems:
  - VDP appliance
  - Data Domain system
  - Backup and restore clients
  - vCenter server
  - vSphere hosts
- Use Hosts files to resolve hostnames to non-routable IP addresses.
- Do not create secondary hostnames to associate with alternate or local IP interfaces.

## NTP Configuration

Configure the VDP appliance, vCenter server, vSphere hosts, and Data Domain systems to use the same Network Time Protocol (NTP) Server.

## Licensing

Ensure that the environment meets the licensing requirements in Table 9-12.

**Table 9-12.** Licensing requirements

| Product | Licensing requirement |
|---|---|
| VDP appliance | The VDP appliance requires a valid vSphere host license (minimum of Essentials Plus). |
| Data Domain system | The DD Boost license must be installed on the Data Domain system. |

## Port Usage and Firewall Requirements

To enable communication between the VDP appliance and the Data Domain systems, review and implement the port usage and firewall requirements that “VDP Port Usage” on page 191 describes.

## Capacity

Carefully assess your backup storage needs when evaluating how much data to store on the Data Domain system and the VDP appliance. Include estimates from data that is sent to the Data Domain system from any other servers.

When the Data Domain system reaches its maximum storage capacity, no further backups to the Data Domain system occur until additional capacity is added or old backups are deleted or expired.

“Monitoring Data Domain from the VDP Appliance” on page 94 provides details on how to monitor capacity.

## Data Domain System Streams

Each Data Domain system has a soft limit to the maximum number of connection and data streams that can be sustained simultaneously while maintaining performance. The number of streams varies depending on the Data Domain system model. For example, the DD990 can support 540 backup streams, while the DD620 can support 20 backup streams.

### Changing the Max Streams Value

As a default, the VDP appliance is configured to use a max streams value of 16.

If you need to modify the number of streams associated to a Data Domain system from the VDP appliance, perform the following steps. The changes applied using these steps will only take effect during subsequent edits of or the addition of a Data Domain system to the VDP appliance.

1. Go to the command line of the VDP appliance (either with SSH / Putty or Terminal of the appliance) and type the following command:

   **cd /usr/local/vdr/etc/**

2. Use a text editor to edit the `vdp-options.properties` file.
3. Insert the field `com.vmware.vdp.option.datadomain.maxstreamsoverride=`*`num`* where *num* is the max number of streams for the Data Domain system.
4. Save the modified file.
5. Add or edit a Data Domain system. Allow five minutes for the appropriate process to run.

The `ddrmaint read-ddr-info` file should now contain a "max-streams" attribute with the value you configured.

## Existing Backup Products in Use with Data Domain

Data Domain systems can use other third-party backup and archiving software. The VDP appliance does not assume it has sole ownership of the Data Domain system. Ensure that proper sizing is evaluated if the system is shared with other software products. The VDP appliance does not use the native Data Domain system snapshot and replication features.

Replication occurs through the DD Boost SDK library by using copying and cloning. However, other third party products may make use of the native Data Domain system snapshot and replication features. In this case, a snapshot of an entire Data Domain system or a replication of an entire Data Domain system includes the VDP appliance data.

# Preparing the Data Domain system for VDP Integration

To support VDP and Data Domain system integration, ensure the environment meets the Data Domain system requirements listed in Table 9-13.

**Table 9-13.** Data Domain system requirements

| Data Domain feature or specification | Requirement for use with the VDP appliance |
|---|---|
| Data Domain Operating System (DD OS) | VDP integration requires DD OS 5.4.0.8 or DD OS 5.5.x or later. |
| DD Boost | VDP integration requires DD Boost 2.6.x DD Boost software enables backup servers to communicate with storage systems without the need for Data Domain systems to emulate tape.<br>There are two components to DD Boost: one component that runs on the backup server and another that runs on the Data Domain system.<br>The component that runs on the backup server (DD Boost libraries) is integrated into the VDP appliance. DD Boost software is an optional product that requires a license to operate on the Data Domain system. |

**Table 9-13.** Data Domain system requirements

| | |
|---|---|
| Data Domain device type | The VDP appliance supports any Data Domain system that is running the supported DD OS version. |
| DD Boost user account | The DD Boost library uses a unique login account name created on the Data Domain system. This account name is known as the DD Boost account. Only one DD Boost account exists per Data Domain system.<br>If the account is renamed and/or the password is changed, these changes must be immediately updated on the VDP appliance by editing the Data Domain configuration options. Failure to update the DD Boost account information could potentially yield integrity check errors and/or backup/restore problems. The DD Boost account must have administrator privileges. |

Before you can add a Data Domain system to the VDP configuration, prepare the Data Domain system by enabling DD Boost and creating a DD Boost user account for the VDP appliance to use to access the Data Domain system for backups and restores (and replication, if applicable).

To prepare the Data Domain system:

1 Disable DD Boost on the Data Domain system by logging into the CLI as an administrative user and typing the following command:

**ddboost disable**

2 Create a DD Boost account and password:

a Create the user account with admin privileges by typing the following command:

**user add USER role admin**

where USER is the username for the new account.

b Set the new account as the DD Boost user by typing the following command:

**ddboost set user-name USER**

where USER is the username for the account.

3 Enable DD Boost to allow the changes to take effect by typing the following command:

**ddboost enable**

IMPORTANT If you change the DD Boost account name or password, make sure to edit the Data Domain system configuration in the VDP configuration utility. Otherwise all backups, restores, and maintenance activities fail.

## Adding a Data Domain System

A Data Domain system performs deduplication through DD OS software. When you select a Data Domain system as the backup target, backup data is transferred to the Data Domain system. Only one Data Domain system can be configured.

### Procedure

1 Log in to the VDP configuration utility by using the following URL:

**https:**<*IP_address_VDP_appliance*>**:8543/vdp-configure/**

2 Click the **Storage** tab.

The storage summary displays statistics about the total usable storage and available capacity for the Data Domain system and for each datastore.

3 From the **Action** list, select **Add Data Domain**.

The **Host Configuration** dialog box appears.

4 Specify Data Domain system information:

   a In the **Data Domain FQDN or IP** box, enter the FQDN or IP address of the Data Domain system to add.

**NOTE** Do not use an IP address or a secondary hostname that associates with alternative or local IP interfaces. It may limit the ability of the VDP appliance to route optimized deduplication traffic.

   b In the **DDBoost User Name** box, type the name of the DD Boost account for VDP to use to access the Data Domain system for backups, restores, and replication.

   c In the **Password** box, type the password for the account that VDP should use to access the Data Domain system for backups, restores, and replication.

   d In the **Confirm Password** box, type the password again to verify it.

   e Click the **Enable Checkpoint Copy** checkbox to enable checkpoint backup support, which allows VDP checkpoints to be stored on a Data Domain system (using DD OS 5.3 or later). These checkpoints are then used if disaster recovery is required.

5 To configure SNMP, click **Next**.

   The SNMP dialog box appears. The SNMP options to configure for VDP and Data Domain system integration include the following:

   - The **Getter/Setter Port Number** text box lists the port on the Data Domain system from which to receive and on which to set SNMP objects. The default value is 161.
   - The **SNMP Community String** text box lists the community string VDP uses for read-only access to the Data Domain system.
   - The **Trap Port Number** text box lists the trap port. The default value is 163.

6 Click **Next**.

   The Ready to Complete dialog box appears.

7 Click **Add** to save your Data Domain configuration.

   A successful Add Data Domain operation causes the following UI changes to occur:

   - The system creates a new checkpoint which takes approximately ten minutes.
   - Data Domain information appears on the VDP appliance in the following locations:
     - **Backup** tab: The Data Domain system is available as the storage target in the Create a new backup job wizard.
     - **Restore** tab: Displays the Data Domain in the Name column of the Restore backup wizard.
     - **Reports** tab: Provides backup status reports for the Data Domain system.
     - **Storage summary**: Displays statistics about the total usable storage and available capacity for the Data Domain. Refer to "Viewing the Storage Configuration" on page 76 for details.
     - **Email reporting**: Displays a summary of the Data Domain configuration.

**NOTE** When you add a Data Domain system to the VDP configuration, the VDP appliance creates an MTree on the Data Domain system for the VDP appliance. The MTree refers to the directory created within the DD Boost path. Data Domain systems support a maximum of 100 MTrees. If you reach the limit, then you cannot add the Data Domain system to the VDP configuration.

## Editing the Data Domain System

1 Log in to the VDP configuration utility by using the following URL:

**https:**<*IP_address_VDP_appliance*>**:8543/vdp-configure/**

2 Click the **Storage** tab.

The storage summary displays statistics about the total usable storage and available capacity for the Data Domain and for each datastore.

3 From the **Action** list, select **Edit Data Domain**.

The **Host Configuration** dialog box appears.

4 Edit the settings for the Data Domain system as necessary. "Adding a Data Domain System" on page 85 provides details on each setting in the dialog box.

5 Click **Next**.

6 After the edits are complete, click **Finish**.

**NOTE** If you edit the Data Domain system hostname, the DD Boost username, or the DD Boost password, the system automatically creates a new checkpoint, which takes approximately ten minutes. "Rolling Back an Appliance" on page 43 provides information about checkpoint.

**NOTE** If you perform a rollback to a checkpoint with the outdated Data Domain system name or DD Boost information, then the rollback fails.

## Deleting the Data Domain System from the VDP Appliance

Before you delete the Data Domain system from the VDP appliance, note the following:

- You must delete all restore points stored on the Data Domain by using the VDP plug-in on the vSphere web client.
- No backup jobs can exist with the Data Domain system. If any backup jobs exist where the Data Domain system is configured as the destination, you must either edit the backup jobs to set a new destination or you must delete the backup jobs.
- After the restore points are checked and the backup jobs are verified, it is best practice to run an integrity check from the **Configuration** tab of the VDP appliance.
- Remove the Data Domain system from VDP using the command line interface. See detailed instructions below.
- After the Data Domain system is deleted, run another integrity check from the VDP user interface to verify the Data Domain system is invalid.

**NOTE** VMware Knowledge Base Article 2063806 provides information about deleting a Data Domain system. This is an internal article. Contact Technical Support for assistance.

### Procedure

1 Before you delete the Data Domain system, delete all restore points stored on the Data Domain system. Use the vSphere web client to delete restore points:

a Go to the **Restore** tab of the VDP plug-in.

b Select the **Manual restore** tab on the navigation bar.

c For clients that have been backed up to the Data Domain system, remove all restore points where the location shows that they are stored on the Data Domain server.

2 Ensure that no backup jobs use the Data Domain system as the destination. If any backup jobs exist with the Data Domain system as a destination, you must edit the backup job to set a new destination or the delete the backup job.

3 After the restore points have been checked and the backup jobs have been verified, as a best practice run an integrity check from the configuration tab of the VDP plug-in.

4 Once the integrity check and validation of the integrity check complete, remove the Data Domain system from the VDP appliance. Use the command line.

   a SSH or Putty into the VDP appliance.

   b Run the `status.dpn` command and verify that the `Last checkpoint` and `Last hfscheck` have completed. If they have not, repeat this step until they show they have completed.

   c Run the `mccli server show-prop` command. This command displays results similar to following output:

| Attribute | Value |
|---|---|
| State | Full Access |
| Active sessions | 0 |
| Total capacity | 575.9 GB |
| Capacity used | 0 bytes |
| Server utilization | 0.0% |
| Bytes protected | 0 bytes |
| Bytes protected quota | Not configured |
| License expiration | Never |
| Time since Server initialization | 1 days 20h:58m |
| Last checkpoint | 2014-10-10 09:03:48 MDT |
| Last validated checkpoint | 2014-10-09 09:02:16 MDT |
| System Name | gs-pod187.test.domain |
| System ID | 1381255529@00:50:56:86:46:10 |
| HFSAddr | gs-pod187.test.domain |
| HFSPort | 27000 |

The System ID contains a number, then an @ sign, and then the MAC address of the VDP appliance. Note the number before the @ sign. The Data Domain system refers to this number as the DPN ID.

5 Run the following command:

**ddrmaint has-backups –dpnid=***num* **–ddr-server=***DDRSERVER* **| grep ‘ hasbackups’**

where *`num`* is the DPN ID you noted in Step c and *`DDRSERVER`* is either the hostname or the IP address of the DDR server. Note that there is a space in the `grep` command between the single quote and the word hasbackups.

This command displays either of the following results:

`hasbackups=”true”`

or

hasbackups=”false”

6 If the information returned is `hasbackups='true'`, then check if Step 1 and Step 2 need to be repeated. After you repeat Step 1 and Step 2 (or verify that the steps have been completed), repeat Step 5.

7 If Step 5 still shows `'hasbackups=true'` continue to Step a. Otherwise, you need to run Step 11.

If you have attempted to remove backups from the Data Domain system by using the VDP UI and the Data Domain system still indicates that the backup data is present, then you must mount the Data Domain data partition to a Linux VM to clear out the data directory. If a Linux VM is unavailable, you can use the VDP appliance for the next steps.

By default, all of the data for backups on a Data Domain system is stored under a single logical storage unit (LSU). The LSU for the VDP is named `Avamar-<DPNID>` and is located under `/data/col1`.

If you are unable to access the file system from the Data Domain operating system interface, you must grant remote access to the LSU. To do so, you must access the Data Domain system remotely by using the following steps:

a Putty or SSH to the Data Domain system.

b Run the `nfs add /data/col1 <IP of Linux VM>` command.

This command should return the results of `NFS export for "/data/col1" added`.

If the command does not return the expected results, type `nfs help` for a man page on the command. If the command returns the expected results, exit the SSH or Putty session.

c Putty or SSH to the Linux VM used in Step b as the root user.

d Run the `mkdir /mnt/DataDomain01` command.

e Run the `mount <IP of DD>:/data/col1 /mnt/DataDomain01` command.

f Run the `ls -ltr /mnt/DataDomain01/avamar-<DPNID>` command, where `DPNID` is the value noted in Step c. The output should display subdirectories where the VDP backups are stored.

g Run the `rm -rf /mnt/DataDomain01/avamar-<DPNID>/*` command where `DPNID` is the value noted in Step c. This will remove all data from the VDP backups.

8 Repeat Step f to verify that all data was removed.

9 Exit the Linux virtual machine.

10 Putty or SSH to the VDP appliance

11 Run the `mccli dd delete --name=<DD IP or hostname> --force=true` command.

12 After you delete the Data Domain system, run an integrity check again from the VDP UI. The old checkpoints with the Data Domain information is now invalid.

# Backups with VDP and Data Domain

The following topics describe VDP and Data Domain system backups:

- How backups work with VDP and Data Domain system
- Selecting a Data Domain target for backups

## How Backups Work with VDP and Data Domain

During a backup, the VDP appliance sends a backup request to the vCenter server. If the backup request includes the option to use a Data Domain system as the target, backup data is stored on the Data Domain system and metadata is stored on the VDP appliance.

The following topics provide additional details on how backups work with VDP and Data Domain.

## Where Backup Data is Stored

All data for a backup is stored under a single dedicated MTree on a single Data Domain system.

## How VDP Appliance Manages Backup Data

During a backup, the VDP appliance sends the metadata for the backup from the client to the VDP data partitions. This process enables the VDP appliance to manage the backup even though the data is stored on a Data Domain system.

The VDP appliance does not store the original path and file name for a file on the Data Domain system. Instead, the VDP appliance uses unique file names on the Data Domain system.

## Supported Backup Types

Store all backup types (full, differential, incremental) for a client in the same destination (VDP appliance or Data Domain system). Backup types should not be distributed across destinations. For example, do not store the initial full backup of a client in the VDP appliance and subsequent differential backups on the Data Domain system.

## Canceling and Deleting Backups

If you cancel a backup while it is in progress, then the VDP appliance deletes the backup data that was written to the Data Domain system during the next cycle of the VDP appliance garbage collection process.

If you delete a backup in VDP, then the backup is deleted from the Data Domain system during the next cycle of the VDP garbage collection process.

"Deleting a Backup Job" on page 119 provides instructions on how to cancel or delete a backup.

## Selecting a Data Domain Target for Backups

After the VDP appliance and the Data Domain system are integrated, any backup target for the VDP appliance has the option to use the Data Domain storage as the "Destination" in the Create a new backup job work flow, as shown in Figure 9-7.

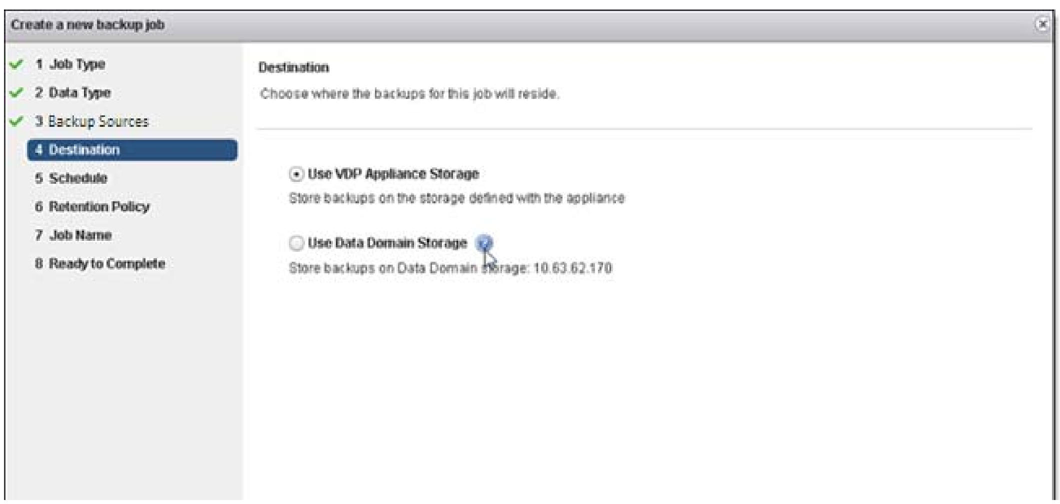



**Figure 9-7.** Create a new backup job wizard - Destination page

You can use the Edit a backup job wizard to change the destination for a backup job. See "Editing a Backup Job" on page 119 for more information on how to edit the backup job.

**NOTE** If the destination of a backup job is modified, then the next backup performed will be a full backup. The new destination will not have the previous full backup data stored with it.

# Replication Control

When a VDP appliance with a Data Domain system attached replicates backup data, replication occurs between the Data Domain systems. However, replication jobs are configured by using the VDP user interface.

You configure and monitor replication on the VDP appliance. The replication activity can also be monitored through the Data Domain system by checking the DD Boost activity. Refer to the *Data Domain Operating Systems Administration Guide* for instructions on how to monitor this activity.

Do not use Data Domain replication functionality to initiate replication of data that was backed up from a VDP appliance. When you use Data Domain replication, the replicated data will not refer to the associated VDP appliance, because the metadata stored on the VDP appliance was not replicated.

### Replication Data Flow

VDP replicates the data directly from one Data Domain system to another. The replication process examines each backup to be replicated, and if it determines that the backup data is stored on a Data Domain system, the process issues a request to replicate the data from the source Data Domain system to the target Data Domain system by using DD Boost. In this instance, the Data Domain systems are responsible for the replication of the data. This is analyzed for each backup being replicated.

### Replication Schedule

The replication of VDP data on a Data Domain system occurs on the VDP replication schedule. You should not schedule replication of data on the Data Domain system separately from the replication of data on the VDP appliance.

### Replication Configuration

To configure a replication when you use a Data Domain system as a backup target for VDP, use the VDP plug-in on the vSphere web client.

Refer to "Replication" on page 139 for more information on configuring VDP replication.

**NOTE** When replicating, if the source is a VDP appliance with a Data Domain system attached to it, the target (whether it be an Avamar server or the VDP appliance) must also have a Data Domain system attached to it.

### Replication Monitoring with VDP

To monitor replication activity with the VDP appliance, including replication activities associated with a Data Domain system, perform the following steps:

1. In the vSphere web client, log in to the vSphere Data Protection plug-in.
2. Click the **Replication** tab.
   - The **Replication** tab displays all replication jobs, and the last run time as well as the next scheduled run time.
   - Selecting a replication job displays the Destination server and the clients included in the replication job details frame.
   - If you check the Replication column, the **Reports** tab displays the replication job and the last replication run time for each protected client.

## Server Maintenance Activity Monitoring

The VDP appliance performs the system maintenance operations for backup data on the Data Domain system, including VDP integrity checks, checkpoints, rollbacks, garbage collection, and secure backup deletion.

The `ddrmaint` utility implements all required operations on the Data Domain system for the VDP appliance. The `ddrmaint` utility logs all maintenance activities on the VDP appliance in the `ddrmaint.log` file. This log file can be located in the `/usr/local/avamar/var/ddrmaintlogs` directory.

When the `ddrmaint.log` file reaches 25 MB in size, the existing log version will roll. The existing `ddrmaint.log` file will be renamed `ddrmaint.log.1` and a new `ddrmaint.log` file will be created. Any previous copies of `ddrmaint.log.X` will have their log counter increased by one as well (`ddrmaint.log.1` will migrate to `ddrmaint.log.2` and so forth).

# Restoring Avamar Checkpoint backups from Data Domain systems

If you have created VDP checkpoint backups on a Data Domain system, you can restore a checkpoint to a new VDP appliance in the event the original VDP appliance fails.

## Assumptions for the Restore Operation

The procedure in "Performing the Checkpoint Restore" on page 92 explains how to perform a checkpoint restore when the following assumptions are true:

- You have a valid checkpoint for a VDP appliance on a Data Domain system target.
- The VDP appliance that failed has been replaced.
- The replacement VDP appliance is a new appliance with no backup data or metadata.
- The size of the replacement VDP appliance is either the same as or larger than the original VDP appliance.
- The replacement VDP appliance must have the same data partition count as the original VDP appliance.
- The replacement VDP appliance must use the same hostname and IP address as the original VDP appliance.

## Performing the Checkpoint Restore

To restore a checkpoint from a Data Domain system to a new VDP appliance:

1. Log in to the VDP appliance as root and from a CLI prompt query the available checkpoints to recover by typing the following command:

   **ddrmaint cp-backup-list --full --ddr-server=***Data_Domain_system* **--ddr-user=***DD_Boost_user_name* **--ddr-password=***DD_Boost_user_password*

   where

   - *Data_Domain_system* is the Data Domain system with the VDP appliance checkpoint backup.
   - *DD_Boost_user_name* is the DD Boost user account used for VDP and Data Domain system integration.
   - *DD_Boost_user_password* is the DD Boost user account password used for VDP and Data Domain system integration.

   The output will be similar to the following example:

   ```
   ================== Checkpoint ==================
   VDP Advanced Name: a4dpe223d
   VDP Advanced MTree/LSU: avamar-1346892530
   Data Domain System Name: griffin-dd10.asl.lab.emc.com
   VDP Advanced Client Path: /MC_SYSTEM/avamar-1346892530
   VDP Advanced Client ID: 8b75468f70dc8ff0fa2e5118cec8ecdddf7fccee
   Checkpoint Name: cp.20140919184604
   Checkpoint Backup Date: 2014-09-19 11:51:12
   Data Partitions: 6
   Attached Data Domain systems: griffin-dd10.asl.lab.emc.com
   ```

2. Restore the backups stored on the VDP appliance from the checkpoint stored on the Data Domain system (requires the Data Domain server name and credentials of the default backup's target Data Domain system) by using the `cprestore` command.

   The `cprestore` command is used for the restore operation. The `cprestore` command completes the following tasks:

   - Creates NFS export on Data Domain system.
   - Mounts Data Domain NFS export on the VDP appliance.
   - Copies the backup files needed on the VDP appliance from the backup checkpoint on the Data Domain system to the corresponding VDP appliance checkpoint directory in each data partition.

- Undoes NFS mount and export.

To restore the backups on the Data Domain system, type the following command on the VDP appliance:

**/usr/local/avamar/bin/cprestore --hfscreatetime=***VDP_ID* **--ddr-server=***Data_Domain_system*
**--ddr-user=***DD_Boost_user_name* **--cptag=***Checkpoint_name*

where:

- *VDP_ID* is determined from the output of Step 1. From the VDP appliance `MTree/LSU:avamar-1346892530` field, the `VDP_ID` is `1346892530`.
- *Data_Domain_system* is the Data Domain system with the VDP appliance checkpoint backup. In the previous checkpoint output example, the value is `griffin-dd10.asl.lab.emc.com`.
- *DD_Boost_user_name* is the DD Boost user account used for VDP and Data Domain system integration. In the previous checkpoint output example, the value is `VDP`.
- *Checkpoint_name* is the checkpoint name. In the previous checkpoint output example, the value is `cp.20140919184604`.

3 Stop the VDP appliance by typing the following command:

**dpnctl stop**

A confirmation message, "Do you wish to shut down the local instance of EMS?" appears. Type **Y**.

4 To initiate a rollback, type the following command:

**dpnctl start --force_rollback**

A message appears that the appliance was shut down. A list of choices also appears.

5 Select option **3, "Select a specific checkpoint to which to roll back"**.

Wait for the rollback to complete. The rollback might take up to one hour, depending on the amount of data present in the VDP appliance. When the rollback is complete, the command prompt returns.

6 Open the user-defined temporary file created during the rollback and verify that the rollback successfully completed without errors.

7 If the internal proxy is enabled, first disable it and then re-enable it by using the VDP configuration utility.

8 If there are any orphan proxies, that is, the external proxies that were deployed on an older VDP appliance, perform the following steps by using the the VDP configuration utility:

a Delete the orphan proxies.

b Change the password of the VDP appliance.

c Add the external proxies.

9 For replication to work after the checkpoint restore, run the following commands:

- `service avagent-replicate stop`
- `service avagent-replicate unregister 127.0.0.1 /MC_SYSTEM`

10 On Avamar Management Console:

a Select **Policy > Clients > MC_system**.

b Right-click the client and select **Edit**.

c In the **Edit Client** dialog box, clear the **Activated** option and click **OK**.

11 On the VDP appliance, run the following commands:

- `service avagent-replicate register 127.0.0.1 /MC_SYSTEM`
- `service avagent-replicate restart - should have the checkbox activated now`

12 On the VDP appliance, create a new checkpoint:

a On the VDP **Configuration** tab, select **Create integrity check** to create a new checkpoint.

b When the checkpoint has been initiated, click OK.

# Monitoring Data Domain from the VDP Appliance

To review high-level information about the capacity of theData Domain system that is attached to a VDP appliance, use either the vSphere web client or the VDP configuration utility.

- On the vSphere web client, open the VDP plug-in and navigate to the **Configuration** tab to view the capacity summary for the Data Domain system.
- In the VDP configuration utility, open the **Storage** tab to view the storage summary for the Data Domain system.

**NOTE** To refresh the data on the page, click the refresh icon beside **Storage Summary**.

When the Data Domain system reaches its capacity limit, you can reclaim space on the device by using the instructions in "Reclaiming Storage on a Full Data Domain System" on page 95.

**Table 9-14.** Capacity thresholds for Data Domain

| Threshold | Value | Behavior |
|---|---|---|
| Capacity warning | 80% | VDP issues a warning event. |
| Capacity error | 95% | Tasks are not generated on the vCenter for backup jobs when Data Domain capacity is more than 95%. |
| Healthcheck limit | 95% | Existing backups are allowed to complete, but new backup activities are suspended. VDP issues error events for Data Domain. |

**NOTE** When the Data Domain system reaches 99% capacity, maintenance operations fail. The best practice recommendation is to limit Data Domain capacity usage to 80%.

## Monitoring by Using the vSphere Web Client

1 On the vSphere web client, open the VDP plug-in.

2 On the **Configuration** tab, the **Storage Summary** section displays the following information:

- Data Domain system FQDN or IP address
- Capacity of the Data Domain system
- Free Space on the Data Domain system
- Used Capacity on the Data Domain system

3 On the left, a color container variant is presented to check the Data Domain capacity:

- When the storage capacity used is less than 80%, the container is green.
- When the storage capacity used is between 80% and 95%, the container is yellow. The following message appears:

  `The Data Domain storage is nearly full.`

- When the storage capacity used is between 95% and 100%, the container is red. The following message appears:

  `The Data Domain storage is nearly full.`

- When the storage capacity used is 100%, the container is red. The following message appears:

  `The Data Domain storage is full.`

### Monitoring by Using the VDP Configuration Utility

1 Log in to the VDP configuration utility by using the following URL:

**https://**<*IP_address_VDP_appliance*>**:8543/vdp-configure/**

2 On the **Storage** tab, the **Storage Summary** section displays the following information:

- DD hostname or IP address
- Total usable storage
- Storage available
- Capacity consumed (in percentage)

**NOTE** To refresh the data on the page, click the refresh icon beside **Storage Summary**.

3 On the left, a color container variant is presented to check the Data Domain capacity:

- When the storage capacity used is less than 80%, the container is green.
- When the storage capacity used is between 80% and 95%, the container is yellow. The following message appears:

  `The Data Domain storage is nearly full.`

- When the storage capacity used is between 95% and 100%, the container is red. The following message appears:

  `The Data Domain storage is nearly full.`

- When the storage capacity used is 100%, the container is red. The following message appears:

  `The Data Domain storage is full.`

## Reclaiming Storage on a Full Data Domain System

If you use all of the storage space on a Data Domain system, the following issues may occur:

- Backups do not succeed and may not start.
- Operations that change information on the Data Domain system fail, including the deletion of checkpoints, active backups, and expired backups during garbage collection. These operations may fail because they involve directory renames, which are not allowed on a full Data Domain system.

To reclaim the used storage on a full Data Domain system, perform the following steps:

1 Determine the source of the data that is using the storage. The data may be from a specific client, a group of clients associated with a specific VDP appliance, or a different backup product that stores data on the Data Domain system.

2 Cancel any backups that are in progress. You must do this from the command line of the VDP appliance.

   a Open an SSH or Putty session to the VDP appliance and run the following commands:

```
su - admin
ssh-agent bash
ssh-add .ssh/dpnid
```

b Run the `mccli activity show` command.

This command returns results similar to the following sample output:

```
admin@gs-pod192:~/>: mccli activity show
0,23000,CLI command completed successfully.
ID              Status    Error Code Start Time        Elapsed    End Time          Type          Progress Bytes New Bytes Client
                          Domain
---------------- --------- ---------- -------------------- ----------- -------------------- ---------------- -------------- --------- ---------
9138660744236309 Running   0 2013-12-09 09:44 MST 00h:27m:25s 2013-12-10 09:44 MST On-Demand Backup 54.3 GB
                          4.2%      Win2008R2-GSClone /10.7.242.175/VirtualMachines
9138660744234709 Completed 0 2013-12-09 09:44 MST 00h:02m:51s 2013-12-09 09:47 MST On-Demand Backup 40.0 GB
                          <0.05%    GermanExchange    /10.7.242.175/VirtualMachines
9138660718256909 Completed 0 2013-12-09 09:39 MST 00h:01m:06s 2013-12-09 09:40 MST On-Demand Backup 40.0 GB
                          <0.05%    GermanExchange    /10.7.242.175/VirtualMachines
9138660744235609 Completed 02013-12-09 09:44 MST 00h:20m:37s 2013-12-09 10:04 MST On-Demand Backup 40.0 GB
                          2.6%      ActiveDirectory   /10.7.242.175/VirtualMachines \
```

To run the command that cancels the running backup jobs, you need to know the appliance password (entered as the 'AppliancePassword' value below). You also need to note the ID of any Running jobs.

c Type the `mccli activity cancel --mcsuserid=`*`MCUser`* `--mcpasswd=`*`AppliancePassword`* `--id=XXXXX` command, where *`AppliancePassword`* is the appliance password and `XXXX` is the ID of the Running job you wish to cancel. This command returns results similar to the following output:

```
admin@gs-pod192:~/>: mccli activity cancel --mcsuserid=MCUser --mcspasswd=Test12345 --id=9138660744236309
0,22205,Backup cancelled via console
AttributValue
--------------------------
activity-id9138660744236309
```

3 Repeat Step c for all jobs in the Running state.

4 Suspend backups and restores. On the VDP appliance, this can be done by disabling the proxies from the command line. Verify with end users that there are no critical backups or restores that must be performed before you run these commands.

a Open an SSH or Putty session to the VDP appliance.

b Run the `service avagent-vmware stop` command.

5 Suspend server maintenance operations on the VDP appliance.

a Log in to the VDP configuration utility by using the following URL:

**https://**<*IP_address_VDP_appliance*>**:8543/vdp-configure**

b If the Maintenance services appear as **Running**, click the **Stop** button

6 On the Data Domain system, manually delete the existing STAGING, DELETED, or cur/DELETED directories for the VDP appliance.

7 Use the Data Domain Enterprise Manager to initiate the Data Domain file system cleaning operation.

This process should free enough space to enable VDP appliance service maintenance operations to complete successfully.

8 Restart server maintenance operations on the VDP appliance.

a Log in to the VDP configuration utility by using the following URL:

**https://**<*IP_address_VDP_appliance*>**:8543/vdp-configure**

b If the Maintenance services appear as **Stopped**, click the **Start** button.

9 Restart the proxies on the VDP appliance so that backups and restores can run.

a Open an SSH or Putty session to the VDP appliance.

b Run the `service avagent-vmware start` command.

# Common Problems and Solutions

This section lists common problems and solutions when you store VDP appliance backups on a Data Domain system.

## Backup Fails if the Data Domain System is Offline

If the Data Domain system is offline when a backup starts, then the backup may take five minutes or more before it fails. The failure occurs because there is a minimum timeout period of five minutes for almost all DD Boost operations.

To resolve the failed backup, set the Data Domain system online, and then retry the backup.

## Rolling Back after Deleting a Data Domain System

If you roll back to a checkpoint after following the procedure for "Deleting the Data Domain System from the VDP Appliance" on page 87, the rollback should bring you to a state where the Data Domain system is removed from the VDP appliance.

To add the Data Domain system back to the VDP appliance, use the VDP Configuration UI. Refer to "Adding a Data Domain System" on page 85 for more information.

If you rolled back to a checkpoint before the deletion of the Data Domain system, then the Data Domain system should be still attached and properly configured. To remove the Data Domain system, follow the procedure for "Deleting the Data Domain System from the VDP Appliance" on page 87.

If a rollback of the appliance leaves you in a different state than these, it is best to contact support to define the proper resolution.

# VDP Disk Expansion 10

This chapter contains the following topics:


- "Prerequisites" on page 100
- "VMFS Heap Size Recommendations" on page 100
- "Performing Disk Expansion" on page 101
- "Performance Analysis" on page 103
- "Disk Expansion with Essentials Plus" on page 103

## Prerequisites

Ensure that your configuration meets the following requirements before disk expansion. Failure to meet these requirements can corrupt VDP and require a restoration from a clone or VDP backup.

- Confirm that the minimum CPU and memory requirements are met for the new configuration:
  - The minimum number of virtual CPUs is 4 for VDP capacity options of 2 TB, 4 TB, 6 TB, and 8 TB.
  - The minimum amount of memory per VM depends on the capacity:

| Capacity size | Required memory |
|---|---|
| 2 TB | 6 GB |
| 4 TB | 8 GB |
| 6 TB | 10 GB |
| 8 TB | 12 GB |

- Confirm that both CPU and Memory Hot-Add are enabled. The CPU and Memory Hot-Add options are disabled by default in the case of an upgraded appliance.

**NOTE**

- If you have an Essentials Plus license, you cannot hot-plug the required memory. You must manually increase the memory assigned to VDP. See "Disk Expansion with Essentials Plus" on page 103 for additional details.
- Confirm that you have available disk space for the expansion. You can check your disk space from the **Storage** tab. See "Viewing the Storage Configuration" on page 76.
- Perform disk expansion during the backup window when no backup jobs or any other VDP tasks are running, such as restores or integrity checks. Ensure that all the VDP services are in a running state before attempting disk expansion. See "Starting and Stopping Services" on page 39 for additional details.
- Confirm that you have administrative privileges in the vCenter. See "User Account Configuration" on page 24 to determine if you have if you have administrative privileges for the vCenter.
- Confirm that VMFS heap size is set to the correct value for the amount of virtual disk space associated with the vSphere host. See "VMFS Heap Size Recommendations" on page 100 for additional details.
- Make a note of the MAC address for the appliance before cloning. The MAC address will be used later in cases where storage has failed.
- Create a clone or backup of the VDP appliance and verify that it is valid before disk expansion.

## VMFS Heap Size Recommendations

VMFS heap size determines the amount of virtual disk space that each vSphere host supports. If you exceed the amount of virtual disk space beyond what is configured for VMFS heap size, the following can occur:

- Virtual machines behave erratically
- Cannot allocate memory error messages appear
- Virtual machines may crash or fail to start

Before disk expansion, confirm that VMFS heap size is configured properly for the new virtual disk capacity. Increasing VMFS heap size increases the amount of memory allocated to the vSphere host kernel and requires a system reboot for the changes to take effect.

VMFS3 and VMFS5 use the same settings and are defined in Table 10-15.

**Table 10-15.** VMFS heap size settings

| Version/build | Default heap amount | Default allowed open VMDK storage per host | Minimum heap amount | Maximum heap amount | Maximum heap value | Maximum open VMDK storage per host |
|---|---|---|---|---|---|---|
| vSphere Host 5.1 Patch 1 (914609) and earlier | 80 MB | 8 TB | N/A | 256 MB | 255 | 25 TB |
| vSphere Host 5.1 Update 1 (1065491) and later | 256 MB | 60 TB | 256 MB | 640 MB | 255 | 60 TB |

**NOTE** vSphere 5.5 and later include a much improved heap eviction process, so that there is no need for the larger heap size that consumes memory. vSphere 5.5 and later support a maximum heap size of 256 MB and enables hosts to access all address space of a 64 TB VMFS.

VMware Knowledge Base Article 1004424 describes the steps to change the VMFS heap size settings.

# Performing Disk Expansion

VDP enables you to expand the datastore capacity by using the **Expand Storage** wizard. Disk expansion enables you to expand the VDP storage to 1 TB, 2 TB, 4 TB, 6 TB, or 8 TB.

**NOTE** You cannot change the provision type from thin provisioning to thick provisioning. The disks inherit the provision type that was assigned to them during initial configuration.

When extending a VMDK which is thick eager-zeroed, the extended part is only thick lazy-zeroed.

VMware Knowledge Base Article 1035823 describes the steps to increase the size of the VMDK and make the VMDK thick eager-zeroed.

## Prerequisite

During the initial configuration, the VDP storage disks are validated.

## Procedure

1 Log in to the VDP configuration utility by using the following URL:

**https://**<*IP_address_VDP_appliance*>**:8543/vdp-configure/**

2 Click the **Storage** tab.

The available datastores appear in the **Storage Summary** section, along with the amount of total usable storage and storage capacity that is available for each datastore.

**NOTE** To refresh the data on the page, click the refresh icon beside **Storage Summary**.

3 From the **Action** list, select **Expand Storage**.

The Expand Storage wizard appears, displaying the current capacity.

4 Determine whether you are expanding the size of the disks or increasing the number of disks.

**NOTE** During expansion to up to 2 TB, the number of disks remains at 3, but the size of the existing disks grows. When disks are expanded, the controls to select the number of disks on each datastore are disabled.

a To expand the size of the disks, select a new capacity from the list. You can expand VDP storage to 1 TB, 2 TB, 4 TB, 6 TB, or 8 TB.

b To add disks, increase the number of disks in the Disks column until the total number of available storage disks are allocated. You can allocate all disks to a single datastore or you can distribute the disks across multiple datastores.

5 Click **Next**.

The Device Allocation dialog box displays the datastores that are known to be available for allocation and the number of disks to be allocated.

A warning message appears if the system detects that a performance analysis has failed, has never been run, or is out of date on one or more of the selected datastores. Based on whether or not you want to run the performance analysis on the selected datastore, do one of the following:

- Click **Yes** to abort the wizard and run the performance analysis on the selected datastore. To run the performance analysis tool, refer to "Performance Analysis" on page 103.
- Click **No** to continue with the disk expansion.

6 Click **Next** to go to the next page of the Expand Storage wizard.

The CPU and Memory dialog box displays the minimum CPU and Memory requirements for the current configuration.

7 Select the number of virtual CPUs for each virtual machine.

The maximum number of virtual CPUs allowed for a VDP virtual machine is 8.

8 Select the amount of memory to be allocated to the VDP virtual machine.

- The minimum amount of memory per VM depends on the capacity:
  - 2 TB capacity requires 6 GB memory
  - 4 TB capacity requires 8 GB memory
  - 6 TB capacity requires 10 GB memory
  - 8 TB capacity requires 12 GB memory
- The maximum amount of memory allowed for a VDP virtual machine is 64 GB.

9 On the Ready to Complete dialog box, click **Finish** to apply the changes.

**NOTE** After a successful storage configuration, the system automatically creates a checkpoint and runs an integrity check.

## Viewing the Storage Configuration

After the storage expansion completes, you can use the **Storage Summary** section on the **Storage** tab to verify the amount of total usable storage and storage capacity that is available for each datastore.

### Known Limitation

Immediately following disk expansion, load balancing occurs and the result is used capacity displays incorrect values. The values appear correctly, however, after the next maintenance window has completed.

### Procedure

1 Log in to the VDP configuration utility by using the following URL:

**https://***<IP_address_VDP_appliance>***:8580/vdp-configure**

2 Click the **Storage** tab.

The **Storage Summary** section displays total usable storage and storage capacity for each datastore.

**NOTE** To refresh the data on the page, click the refresh icon beside **Storage Summary**.

# Performance Analysis

The performance analysis test performs the write, read, and seek performance tests. During the initial configuration, the performance analysis test evaluates the read, write, and seek capabilities of the configured VDP disks. After the initial configuration, the tests measure the performance of the datastores by creating a temporary disk of 256 GB on the datastores.

Possible test results are as follows:

- Unknown
- Failed
- Passed

## Running the Performance Analysis Test

1 Log in to the VDP configuration utility by using the following URL:

**https://**<*IP_address_VDP_appliance*>**:8543/vdp-configure**

2 On the **Storage** tab, click **Performance Analysis**.

The performance analysis test creates a temporary disk on the selected datastore and runs the test on that disk. The temporary disk is automatically removed once the performance analysis completes.

3 Select the datastores on which the performance analysis test will run.

4 Click **Run** to start the performance analysis.

NOTE You can click **Cancel** at any time to stop a performance analysis test.

# Disk Expansion with Essentials Plus

If you have an Essentials Plus license, you cannot enable Memory Hot-Add. You must manually adjust the memory assigned to VDP, as required for the target capacity.

NOTE This limitation is true with vSphere 5.x hosts. You must manually set the minimum required memory for a vSphere host using the memory requirements listed in Table 10-16.

VDP requires the following memory based on capacity size.

**Table 10-16.** Memory requirements for virtual hardware

| Capacity size | Required memory |
|---|---|
| 2 TB | 6 GB |
| 4 TB | 8 GB |
| 6 TB | 10 GB |
| 8 TB | 12 GB |

### Procedure

To perform disk expansion with an Essentials Plus license complete the following steps.

1 From a web browser, access the vSphere web client:

**https://**<*IP_address_vCenter_Server*>**:9443/vsphere-client/**

2 Before the expansion, shut down the VDP appliance by using the **Shut Down Guest OS** action on the virtual machine.

This action automatically performs a clean shutdown of the VDP appliance. If the appliance is powered off without the **Shut Down Guest OS** action, corruption might occur. It can take up to 30 minutes to shut down and restart the VDP appliance. You can monitor the status through the virtual machine console.

3 Increase the memory assigned to the VDP appliance by using the requirements listed in Table 10-16.

   a From a web browser, access the vSphere web client:

      **https://**<IP_address_vCenter_Server>**:9443/vsphere-client/**

   b In the vSphere web client, log in as a user who has privileges to edit hardware settings.

   c Click **vCenter > Hosts and Clusters**.

   d In the tree on the left, click the disclosure arrows until the VDP appliance appears.

   e Right-click the VDP appliance and select **Edit Settings**.

   f Click the **Virtual Hardware** tab.

   g Increase the amount of memory by entering the value in the **Memory** field.

   h Click **OK**.

4 To restart the VDP appliance, right-click the VDP appliance and select **Power On**.

# Using VDP 11

This chapter includes the following topics:


- "Accessing VDP" on page 106
- "Accessing the VDP Appliance by Using CLI" on page 106
- "Understanding the VDP User Interface" on page 107
- "Switching VDP Appliances" on page 108
- "VDP User Interface" on page 108
- "Viewing Information from the Reports Tab" on page 108

# Accessing VDP

Use vSphere web client to access and manage VDP.

**NOTE** You cannot use VDP without a vCenter server. In linked mode, the VDP appliance works only with the vCenter server with which it is associated. If the VDP appliance fails to appear in the vSphere web client, remove your vCenter from linked mode.

### Prerequisites

Before using VDP, you must install and configure the VDP appliance described in "VDP Installation and Configuration" on page 21.

### Procedure

1 From a web browser, access the vSphere web client:

**https://<*IP_address_vCenter_Server*>:9443/vsphere-client/**

2 In the Credentials page, enter an administrative vCenter username and password, and click **Login**. The specified user account must have administrative privileges.

3 In the vSphere web client, select **VDP**.

4 On the **Welcome to VDP** page, select the VDP appliance and click **Connect**.

# Accessing the VDP Appliance by Using CLI

Currently, users can access the VDP appliance command line using the vSphere Client console, SSH, or Putty sessions. With the VDP 5.8 and later releases, the ability to use SSH or Putty to log on to the VDP appliance with the root user has been removed.

To use SSH or Putty to access the appliance, you must log in as the admin user.

With the VDP 5.8 and later releases, the default password for the admin user has also changed. This password only applies when the user logs in to the VDP appliance before completing the configuration of the appliance from the VDP Configuration UI. In this scenario, the admin password is `88RttoTriz!!`

After the configuration of the appliance is complete, the newly-configured appliance password can be used to access the command line as the admin user.

## Useful Commands

The following table lists some of the useful commands:

**Table 11-17.** Useful VDP commands

| Command | Description |
|---|---|
| status.dpn | Displays VDP status. |
| dpnctl status | Displays VDP service status. |
| emwebapp.sh - - test | Tests the VDP Enterprise Manager web application status. |
| capacity.sh | Analyzes space consumption from the last 30 backup jobs, and displays the size of new data and the space that the garbage collection has recovered. |
| cplist | Displays checkpoint status. |
| mccli server show-prop | Displays the VDP appliance properties that is the same as that appears in the vSphere web client. |
| mccli activity show | Displays backup jobs information. Each activity is a backup job from a single virtual machine. |
| mccli dd show-prop | Displays Data Domain systems properties. |

**Table 11-17.** Useful VDP commands

| Command | Description |
|---|---|
| mccli client show --domain=/vCenter FQDN or IP/VirtualMachines --recursive=true | Lists the protected virtual machines. |
| mccli backup show --name=/ vCenter FQDN or IP/VirtualMachines/Virtual Machine name --recursive=true | Lists the backups of a virtual machine. |
| mccli -help | Lists the available command options. |

# Understanding the VDP User Interface

The vSphere web client for vSphere Data Protection is used to configure and manage VDP.

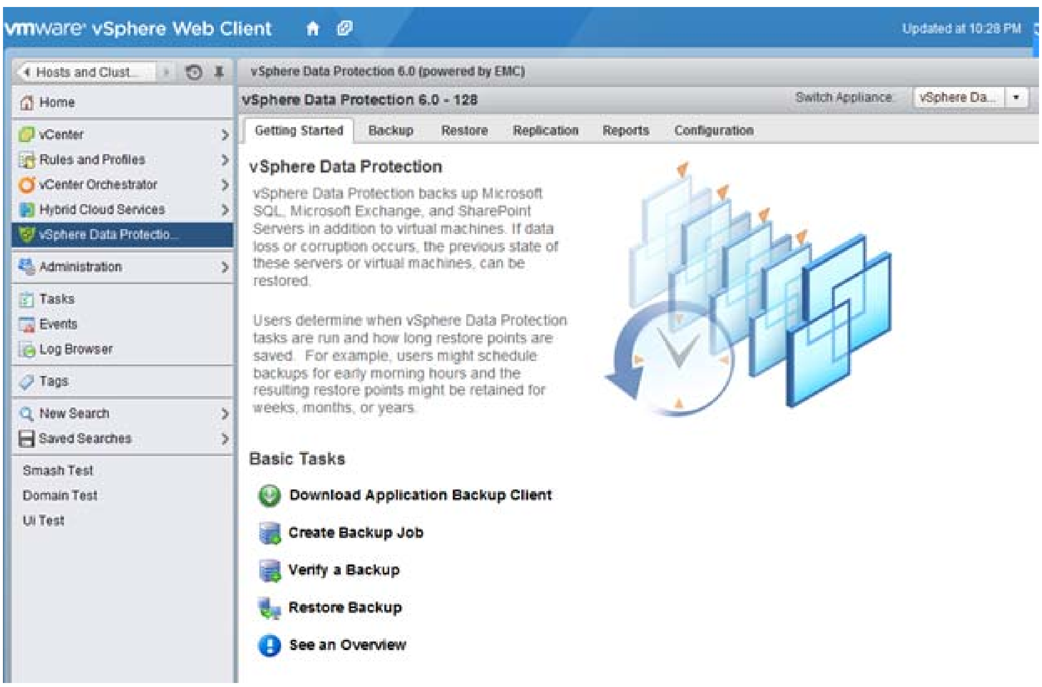



**Figure 11-8.** vSphere Data Protection user interface

The vSphere Data Protection user interface consists of six tabs:

- **Getting Started**: Provides an overview of VDP functionality and quick links to the Download Application Backup Client, Create Backup Job, Verify a Backup, Restore Backup and See an Overview.
- **Backup**: Provides a list of scheduled backup jobs as well as details about each backup job. You can also create and edit backup jobs from this page. This page also provides the ability to run a backup job immediately. "Viewing Information from the Reports Tab" on page 108 provides more information.
- **Restore**: Provides a list of successful backups that you can restore. "Restore Operations" on page 132 provides more information.
- **Replication**: Provides a list of successful backups that you can replicate. "Replication Jobs" on page 140 provides more information.
- **Reports**: Provides backup status reports on the virtual machines on the vCenter server. "Viewing Information from the Reports Tab" on page 108 provides more information.
- **Configuration:** Displays information about how VDP is configured and enables you to edit some of these settings. "Configuring VDP" on page 53 provides more information.

# Switching VDP Appliances

Each vCenter server supports 20 VDP appliances. You can switch appliances by choosing an appliance from the drop-down list to the right of the **Switch Appliance** option.

NOTE The VDP appliances in the drop-down list are sorted alphabetically, and the first item in the list that is displayed on the screen may not match the current appliance. On the VDP screen, the appliance name on the left is the current appliance, and the appliance name in the drop-down list is the first in the list of available appliances.

# VDP User Interface

The VDP user interface includes the following features:

- Guest-level backups of Microsoft Exchange Servers, SQL Servers, and SharePoint Servers.
- Growing and adding disks (disk expansion)
- Backing up to a Data Domain system
- Granular level recovery (GLR)
- Verifying automatic backups

These options are described in "VDP Application Support" on page 162, which also specifies how to perform application-level backups and restores.

# Viewing Information from the Reports Tab

The **Reports** tab displays the following information:

**Appliance Status Information**

- **Appliance status:** The status of the VDP appliance.
- **Integrity check status:** Click the green right arrow to initiate the integrity check. The status value is either Normal or Out of Date.
  - Normal indicates that a successful integrity check has been completed in the past two days.
  - Out of Date indicates that an integrity check has not run or has not completed successfully in the past two days.
- **Used capacity:** A percentage of the total VDP capacity that is occupied by backups.
- **Used DDR capacity:** A percentage of the total capacity that is occupied by the Data Domain system (if applicable.
- **Recent failed backups:** The number of virtual machines that failed to back up in the most recently completed backup job.
- **Recent failed backup verifications**: The number of backup verification jobs that have recently failed.
- **Recent failed replications:** The number of replication jobs that have recently failed.
- **Total VMs protected:** The total number of virtual machines that are protected on the VDP appliance.

NOTE The VDP appliance supports 400 virtual machines. If the maximum numbers of virtual machines are exceeded, an alarm is generated and the **Log** tab displays the error message.

## Refresh

Click the **Refresh** button at any time to update the data in the **Reports** tab.

## Task Failures Tab

The **Task Failures** tab displays details about jobs that have failed within the last 72 hours.

**Failure Time:** The date and time the job failed.

**Reason:** The reason the job failed.

**Client/Source Name:** The client associated with the vCenter.

**Job Name:** The name of the job that failed.

**Job Type:** The type of job that failed; for example, Scheduled Backup or On Demand Backup.

**Next Run Time:** The date and time the job is next scheduled to run.

**View Log**: Click to launch the Job Failure Client Log dialog box, where you can view the client log details. If log files are not available in the Client Log dialog box, one of three error messages may appear:

- Failed to retrieve log file. This message appears in the following circumstances:
  - Management services were recently restarted.
  - Regular log maintenance has removed old log files.
  - Logs may be empty or non-existent.
  - An error may have occurred.
- Failed to retrieve log. Logs are regularly removed 72 hours after the job finished.
- Log file retrieved is empty.

NOTE Not all replication source and replication destination logging is available through the **Job Failures** tab. However, you can retrieve replication logs from the VDP configuration utility by using the regular log bundle. See "Collecting VDP Logs or Diagnostic Information" on page 40 for more information.

### Actions icon

You can perform the following tasks from the Actions icon list, located on the right side of the **Task failures** tab:

- **Rerun Task**: Highlight the failed job and click to rerun the failed job. The Rerun job feature is not enabled for Restore failures.

NOTE To run a backup only for the failed client, select **Backup only out of date source** under **Backup Now** on the **Backup** tab.

- **Export to CSV**: Click to export the current table to a comma-separated values (.CSV) file.
- **Show All Columns**: Hide one or more columns by clicking **X** on the column name, and then click **Show All Columns** to show the hidden columns on the UI.

### Filter

You can filter and customize the job failure details by selecting one of the following criteria. The information appears for job failures that occurred within the last 72 hours.

- **Show All:** Shows all job failure information for the virtual machines. **Show All** is the default.
- **Error:** Filters job failure information by errors messages.
- **Job:** Filters job failure information for a selected job.
- **Client:** Filters job failure information for a given client.

## Job Details Tab

The **Job Details** tab allows you to select the type of job (Backups, Replication, or Backup verification) and display the details for the selected job. Backups is the default job type.

Job details are grouped into three sections:

- **Client Information**
  - **Client Name**: The client associated with the vCenter. Regular VM clients and retired VM clients from the Replicate domain displays the hashed mask value appended to the replicated, recovered, and imported names.
  - **Type**: Displays the type as Image, MS SQL Server, MS SharePoint Server, or MS Exchange Server. Applications (MS SQL Servers, MS SharePoint Servers, or MS Exchange Servers) are supported only on the VDP appliance.
  - **Jobs**: The job name. Multiple job names appear if a virtual machine resides in two different jobs.
- **Last Execution**
  - **Job Name**: The name of the job.
  - **Completion**: The date and time the job completed.
  - **Result**: Whether the job succeeded, failed, succeeded with exceptions, or was canceled.

    If the job succeeded with exceptions, click the **View Log** link to review the logs for information about the exceptions.
- **Next Execution**
  - **Job Name**: The next scheduled job name appears. If a virtual machine resides in two different jobs with different schedules, the next scheduled job name appears.
  - **Scheduled**: The date and time the job is next scheduled to run.

### Actions icon

You can perform the following tasks from the Actions icon list, located on the right side of the **Job Details** tab:

- **Export to CSV**: Click to export the current table to a comma-separated values (.CSV) file.
- **Show All Columns**: Hide one or more columns by clicking **X** on the column name, and then click **Show All Columns** to show the hidden columns on the UI.

### Filter

You can filter and customize the job details by selecting one of the following criteria. The information appears for job failures that occurred within the last 72 hours.

- **Show All:** Shows all job details for the virtual machines. **Show All** is the default.
- **Client:** Filters job failure information by client.
- **Last Execution:** Filters job details for the last executed job.
- **Next Execution:** Filters job details for the next scheduled job.

## Unprotected Clients Tab

- **Client Name**: The name of the unprotected client.
- **IP Address:** The IP address or hostname of the unprotected client.
- **VM Path:** The path where the virtual machine resides.

### Actions icon

You can perform the following task from the Actions icon list, located on the right side of the **Unprotected Clients** tab:

- **Export to CSV**: Click to export the current table to a comma-separated values (.CSV) file.

# Managing Backups 12

This chapter includes the following topics:


- "Backup Jobs" on page 114
- "Choosing the Virtual Machines" on page 114
- "Specifying the Backup Schedule" on page 115
- "Setting the Retention Policy" on page 115
- "Creating a Full Image Backup Job" on page 116
- "Creating a Backup Job on Individual Disks" on page 117
- "Viewing Status and Backup Job Details" on page 119
- "Editing a Backup Job" on page 119
- "Cloning a Backup Job" on page 119
- "Locking and Unlocking a Backup" on page 121
- "Migrating Backup Jobs from VDP to Avamar" on page 121

# Backup Jobs

Backup jobs consist of a set of one or more virtual machines that are associated with a backup schedule and specific retention policies. Backup jobs are created and edited on the **Backup** tab using the Create a new backup job wizard.

### Limitations

- VDP will not back up the following specialized virtual machines:
  - VDP appliances
  - vSphere Storage appliances (VSA)
  - VMware Data Recovery (VDR) appliances
  - Templates
  - Secondary fault tolerant nodes
  - Proxies
  - Avamar Virtual Edition (AVE) servers

  The wizard allows you to select these VMs. When when you click **Finish** to complete the wizard, you receive a warning that these specialized virtual machines were not added to the job.
- Virtual machines that contain special characters in their names cannot be added to any backup job. The following characters cannot be used: ~!@$^%(){}[]|,`;#\/:*?<>'"&. In addition, diacritical characters cannot be used (for example: â, é, ì, ü, and ñ).
- Using snapshots to perform a backup on VMs configured with bus sharing is not supported. If you require SCSI bus sharing, refer to the following Knowledge Base article: http://kb.vmware.com/kb/1006392.

### Prerequisites

- VDP is installed and configured on your vCenter server.
- You are logged in to the vSphere web client and connected to the VDP appliance.

# Choosing the Virtual Machines

You can specify collections of virtual machines, such as all virtual machines in a datacenter or select individual virtual machines. If an entire resource pool, host, datacenter, or folder is selected, any new virtual machines added to that container are included in subsequent backups. If a single virtual machine is selected, any disk subsequently added to the virtual machine is included in the backup. If a virtual machine is moved from the selected container to another container that was not originally selected, it is no longer part of the backup.

You can manually select a virtual machine to be backed up, which ensures that virtual machine is backed up, even if it is moved.

## Identifying retired virtual machines

The following conditions can cause a VM client to retire, and therefore make the client unavailable for future backup jobs:

- **Host**: Removed from inventory
- **Virtual machine**: Deleted from disk
- **Virtual machine**: Removed from inventory

The following conditions exist when the VM client is not retired, and the child VM remains in the inventory:

- **Resource pool**: Removed from inventory
- **vApp**: Removed from inventory
- **Host**: Disconnected
- **Host**: Enters maintenance mode
- **Host**: Shuts down

## Specifying the Backup Schedule

On the Schedule page of the Create new backup job wizard (Step 7) of the procedure below, you can specify the time intervals to back up the virtual machines in your backup job. Backups occur as near to the startup of the backup window as possible. The available time intervals are daily, weekly, or monthly.

## Setting the Retention Policy

Retention policies for the VDP appliance are set individually per backup job. As each restore point is created from a backup job, it retains the retention policy at the time it was created. If a backup job's retention policy is modified, the new policy only affects newly-created restore points. Previously-created restore points retain the previous retention policy.

On the Retention Policy page of the Create Backup Job wizard (Step 8), you specify the retention period for backups.

The first three options, **Forever**, **for**, and **until**, apply to all the backups of all the virtual machines in the group equally. The fourth option, **this Schedule** or **Custom Retention Schedule**, applies only to backups that are internally assigned special Daily, Weekly, Monthly, or Yearly tags.

NOTE The **this Schedule** default is 60 days. The **Custom Retention** default is Never.

The first backup of the day receives a Daily tag. If this backup is also the first backup of the week, the month, and the year, the backup also receives the Weekly, Monthly, and Yearly tags.The time intervals specified by the **this Schedule** or **Custom Retention Schedule** options only apply to backups with the internal tags. Table 12-18 describes Retention Policy options.

**Table 12-18.** Retention Policy Options

| Option | Description |
|---|---|
| **Forever** | Sets backup job to never expire. All backups for the virtual machines in this backup job are never deleted. |
| **for** | Sets a specific number of days, weeks, months, or years for the backup job. All backups for the virtual machines in this backup job are stored until the specified time interval has elapsed from their creation date. For example, if you set a retention policy to 30 days for a backup job, each job that runs has an expiration date of 30 days in the future. |
| **until** | Sets a specific expiration date. All backups for the virtual machines in this backup job are deleted on the date specified in the **until** field. |
| **this Schedule** or **Custom Retention Schedule** | Sets the retention time intervals for backups that are assigned internal tags of Daily, Weekly, Monthly, or Yearly.<br>Backups can have more than one of these internal tags. The tag with the longest time interval has precedence. For example, if you set backups with a Weekly tag to be retained for 8 weeks, and backups with the Monthly tag to be retained for 1 month, then backups with both the Weekly and Monthly tags would be retained for 8 weeks. |

# Creating a Full Image Backup Job

### Prerequisites

- VDP is installed and configured on your vCenter server.
- The disks are supported by VDP. VDP does not support the following virtual hardware disk types:
  - Independent
  - RDM Independent - Virtual Compatibility Mode
  - RDM Physical Compatibility Mode

### Procedure

1 From a web browser, access VDP.

2 Click the **Backup** tab.

The **Backup** tab displays a list of the backup jobs that have been created. The columns in the list are described in the following table:

**Table 12-19.** Column descriptions on **Backup** tab

| Column | Description |
|---|---|
| Name | The name of the backup job. |
| State | Whether the backup job is enabled or disabled. Disabled backup jobs will not run. |
| Type | The type of backup, such as Application or Image. |
| Last Start Time | The last time the backup job was started. |
| Duration | How long it took to complete this job the last time it ran. |
| Next Run Time | When the backup job is scheduled to run again |
| Success Count | The number of virtual machines that were backed up successfully the last time the backup job ran. This number is updated after each backup job |
| Failure Count | The number of virtual machines that did not back up successfully the last time the backup job ran. This number is updated after each backup job. |

3 From the **Backup Job Actions** menu, select **New** to run the Create new backup job wizard.

You can also run the Create a new backup job wizard from the **Getting Started** tab. To do this, click **Create Backup Job** under Basic Tasks.

4 On the **Job Type** page, select **Guest Images**, and then click **Next**.

The **Applications** option enables you to create backup jobs on Microsoft Exchange Servers, Microsoft SQL Servers, and Microsoft SharePoint Servers. "VDP Application Support" on page 161 provides details.

5 On the **Data Type** page, click **Next**.

NOTE Selecting the **Fallback to non-quiescence failure** option enables backups to be performed as non-quiescent in case of a quiescent failure, and marks the restore points as crash-consistent.

6 On the **Backup Sources** page, click the disclosure arrows to progressively disclose the VMs. Select the checkboxes next to the items to add to the backup job, and then click **Next**.

NOTE If a Data Domain system is configured as the backup target, there is an additional step to configure the target destination. Select either **local storage** or **Data Domain storage**.

7 On the **Schedule** page, select the schedule for the backup job and click **Next**.

8 On the **Retention Policy** page, select a retention period and click **Next**. The retention periods from which you can choose are described as follows.

- **Forever**: All backups for the virtual machines in this backup job will never expire.
- **for**: All backups for the virtual machines in this backup job will expire after the specified time interval has elapsed from their creation date. The time interval can be specified in days, weeks, months, or years.
- **until**: All backups for the virtual machines in this backup job will expire on the date specified.
- **this Schedule**: Specifies the retention time intervals for backups that are assigned internal tags. When you perform scheduled daily backups on a regular basis, some of the backups are automatically assigned one of the following retention types:
  - **Daily**: The first successful scheduled backup each day.
  - **Weekly**: The first successful scheduled backup each week.
  - **Monthly**: The first successful scheduled backup each month.
  - **Yearly**: The first successful scheduled backup each year.

The **this Schedule** default for on-demand backup jobs is 60 days.

For the purpose of assigning retention types, each day begins at 00:00:01 GMT, each week begins on Sunday, each month begins on the first calendar day of that month, and each year begins on January 1.

As backups may have more than one of these internal tags, the tag with the longest time interval has precedence. For example, if you were to set backups with a Weekly tag to be retained for 8 weeks, and backups with the Monthly tag to be retained for 1 month, then backups that were assigned both the Weekly and Monthly tags would be retained for 8 weeks.

**CAUTION** Upon entering a new maintenance period following the expiration of a backup, the VDP appliance removes its reference to the backup data and thereafter you cannot restore the expired backup. The VDP appliance determines if the backup data is being used by any other restore point and, if the system determines that the data is not being used, the data is removed and the disk capacity frees up.

9 On the **Name** page, type a name for the backup job and click **Next**.

The backup job name must be unique and can be up to 255 characters long. The following characters cannot be used in the backup job name: ~!@$^%(){}[]|,`;#\/:*?<>'"&. In addition, diacritical characters cannot be used (for example: â, é, ì, ü, and ñ).

10 On the **Ready to Complete** page, review the summary information for the backup job and click **Finish**.

An information dialog box will confirm the backup job was created successfully. The backup operation can take several minutes.

11 Click **OK**.

The newly created backup job is now listed on the **Backup** tab.

## Creating a Backup Job on Individual Disks

A full image backup job aggregates all disks in the entire virtual machine into a single image backup. Individual disk backup jobs allow you to select only the disks you need. This capability allows you to filter based on certain configuration criteria, such as by operating system or by retention policy.

## Unsupported disk types

When planning for individual disk backups, make sure the disks are supported by VDP. Currently, VDP does not support the following virtual hardware disk types:

- Independent
- RDM Independent - Virtual Compatibility Mode
- RDM Physical Compatibility Mode
- Virtual disks attached to the SCSI controller with bus-sharing enabled

**NOTE** If a virtual machine contains a VMDK that is not supported, the VMDK is grayed out and the checkbox is unavailable.

For more information about backing up unsupported disk types, refer to "Single VMDK Backup and Restore" on page 16.

## Limitation

To perform a backup of a single VMDK on the virtual machine that includes multiple VMDKs, the datastore must have enough space for snapshots of all of the VMDKs on the virtual machine. Even though the backup job is meant to back up a single VMDK, the backup process performs a snapshot of all of the VMDKs on the virtual machine.

### Prerequisites

The VDP appliance is installed and configured on your vCenter server.

### Procedure

1. From a web browser, access VDP.
2. Click the **Backup** tab and, from **Backup Job Actions**, click **New** to launch the Create a new backup job wizard.

**NOTE** You can also launch the Create a new backup job wizard from the **Getting Started** tab. To do this, click **Create Backup Job** under Basic Tasks.

3. To back up individual virtual machine disks, select **Individual Disks** as the data type, and then click **Next**.

   The Virtual Machines page displays an inventory tree. This tree contains all the objects and virtual machines in the vCenter server.

   Click on the disclosure arrow to progressively disclose the contents of the tree. Select the checkboxes next to the items to add to the backup job, and then click **Next**.

4. On the Schedule page, select the schedule for the job and click **Next**.
5. On the Retention Policy page, accept the default retention policy or specify an alternate retention policy and click **Next**.
6. On the Name page, enter a backup job name and click **Next**.

   The backup job name must be unique and can be up to 255 characters long. The following characters cannot be used in the backup job name: ~!@$^%(){}[]|,`;#\/:*?<>'"&. In addition, diacritical characters cannot be used (for example: â, é, ì, ü, and ñ).

7. On the Ready to Complete page, review the summary information for the backup job, and then click **Finish**.

   An information dialog box will confirm the backup job was created successfully. The backup operation can take several minutes.

8. Click **OK**.

The newly-created backup job is now listed on the **Backup** tab.

### Migration on Individual Disks

VMware Storage VMotion (SVMotion™) is a component of VMware vSphere that provides an intuitive interface for live migration of vmdk files with no downtime or disruption in service. You can find complete information about migrating with vMotion and SVMotion at the VMware vSphere Documentation Center web site:

http://pubs.vmware.com/vsphere-51/index.jsp.

Users have two options when migrating a virtual machine from one datastore to another:

- Migrate the full virtual machine all at once to another datastore.
- Migrate individual disks to another datastore, where disks for a single virtual machine may reside on a different datastore.

When a full virtual machine is migrated, the VDP appliance updates the backup jobs with the new locations of the protected VMDKs.

From the VDP 6.0 and later releases, when you migrate individual disks (vmdk files) from one datastore to another, VDP recognizes the change in datastore and appropriately backs up vmdk during the next backup.

## Viewing Status and Backup Job Details

The **Backup** tab displays a list of backup jobs that have been created with VDP. By clicking on a backup job, you can see the details of the job on the Backup Job Details pane:

- **Name**: The name of the backup job.
- **Status**: Whether the backup job is enabled or disabled.
- **Sources**: A list of the virtual machines in the backup job.
- **Out of Date**: A list of all the virtual machines that failed to back up the last time the job ran.

## Editing a Backup Job

After you create a backup job, you can edit the job by highlighting the backup job and selecting **Backup Job Actions** > **Edit**.

## Cloning a Backup Job

After you create a backup job, you can use the job as a template for creating a different job by highlighting the backup job and selecting **Backup Job Actions** > **Clone**.

Performing the clone action launches the Cloning backup job wizard and uses information from the original job to automatically fill in the first three pages of the wizard (Virtual Machines, Schedule, and Retention Policy). The cloned job requires a unique name. Except for the data type (because an image backup cannot be changed to an individual disk backup and vice versa), any of the settings that were copied from the original job can be modified.

**NOTE** You can clone full image backups and individual disk backups.

### Deleting a Backup Job

You can delete the job by highlighting the backup job and selecting **Backup Job Actions** > **Delete**.

**NOTE** When using **Delete** on the **Backup** tab you are only deleting the job. Any backups previously made by the job are still retained by VDP in accordance with the retention policy of the job. To delete backups, use **Delete** on the **Restore** tab.

## Enabling or Disabling a Backup Job

If you want to temporarily stop a backup job from running in the future, you can disable it. You can edit and delete disabled backup jobs. VDP will not run a disabled job until it has been enabled.

You can enable or disable backup jobs by highlighting the backup job and selecting **Backup Job Actions** > **Enable/Disable**.

## Running Existing Backup Jobs Immediately

You can run backup jobs immediately by using one of the following methods:

- Choosing to backup up a protected virtual machine
- Choosing to run an existing backup job

### Immediately Backing Up a Protected Virtual Machine

1 Select the protected virtual machine you want to immediately back up by using one of the following options:

- Right-click the virtual machine in an inventory tree and select **All VDP Actions** > **Backup Now**. The virtual machine must belong to a backup job for this selection to appear.
- Click the virtual machine in an inventory tree, and then click the **Actions** button. Select **All VDP Actions** > **Backup Now**. The virtual machine must belong to a backup job for this selection to appear.
- Click the virtual machine (on the **Reports** tab), click the Actions icon, and then select **Backup Now**.

The Backup Now dialog appears.

2 Select the VDP appliance and the backup job, and then click **OK**.

An information dialog appears telling you the backup job has been initiated.

3 Click **OK**.

VDP starts the backup job.

### Immediately Running a Backup Job

1 From the **Backup** tab in the VDP user interface, click the job you want to run immediately.

Multiple selections are allowed on the **Backup** tab by using Ctrl- or Shift-click. Hold down the Ctrl key and click multiple, specific backup jobs. Hold down the Shift key and click a range of backup jobs.

2 Click **Backup Now**.

A drop-down selection appears with the following options:

- **Backup all Sources**: Backs up all the virtual machines in the backup job.
- **Backup only out of date sources**: Backs up only the virtual machines that did not back up successfully the last time the backup job ran.

3 Click the sources you want to back up immediately.

4 Click **OK** when you see the message that the backup has been requested.

VDP starts the backup job.

The **Backup Now** option immediately initiates backup jobs if VDP is in the backup window or the maintenance window.

# Locking and Unlocking a Backup

During maintenance periods, VDP examines the backups in the appliance and evaluates whether the backup retention period has expired. If it has expired, VDP removes the expired backup from the appliance. However, if you want to prevent VDP from deleting a backup, you can lock it. VDP will not evaluate the retention period on that backup again, until it is unlocked.

**NOTE** Data in the VDP database drives the locked status. The VDP database is cleared when disks are imported (see "Attaching Existing VDP Disks" on page 74). When disks are imported, the original expiration date for locked backups is reassigned to "never," and therefore, those disks cannot be unlocked.

**NOTE** You cannot lock individual disk backups. You can only lock full image backups.

### Prerequisites

- VDP is installed and configured on your vCenter server.
- You are logged in to the vSphere web client and connected to the VDP appliance.

### Procedure

1. From a web browser, access VDP.
2. From the **Restore** tab in the VDP user interface, click the disclosure arrow associated with the backups shown in the table to locate the backup you want to lock.
3. Select the checkbox next to the backup you want to lock.
4. Click the **Lock/Unlock** icon.

   Locking a backup overlays a lock icon on the backup icon ( ). The backup is now locked.
5. To unlock a backup, select the **Lock/Unlock** icon again.

   The lock overlay is cleared and VDP evaluates the retention date of the backup during the next maintenance period.

# Migrating Backup Jobs from VDP to Avamar

VDP 6.1 enables you to migrate backup jobs and backup verification jobs to Avamar. The migration makes the backup group and the VM backup validation group visible and fully executable on the target Avamar server.

After the migration completes:

- There will be a new records of Groups, Clients, and related records such as, Datasets, Retentions, Schedules, and so on, on the target Avamar server.
- Avamar can use the same policies to protect the same VM clients as VDP did before the migration.

## Guidelines

- VDP migration is a one-time process. Once executed, VDP disables further migrations.
- The migration process migrates only the details of backup jobs and backup verification jobs.
- In case of application backup jobs, reregister the application clients on the Avamar server instead of the VDP server to properly execute the newly migrated jobs.

## Recommendations

- You replicate all the existing backups to the target Avamar server before you perform migration.
- If you migrate backup jobs from VDP 6.1.3 to Avamar 7.3:

  a. Before you start migration, stop the avfirewall service on both the Avamar appliance and VDP by running the following command:

     service avfirewall stop

  b. After the migration completes, start the avfirewall service on both the Avamar appliance and VDP by running the following command:

     service avfirewall start

## Prerequisites

- You have registered the source vCenter client to the target Avamar server. Otherwise, you cannot perform the migration, and must manually add the source vCenter as a new VMware vCenter client on the target Avamar server.
- You have deployed and configured VMware Image Proxy on the target to work with the source vCenter.
- You have ensured that the target Avamar server contains at least one Data Domain system if the source VDP appliance is configured with a Data Domain system.

## Procedure

1. Open a web browser and type the following URL in the address bar:

   **https://**<IP_address_of_VDP_Appliance>**:8543/vdp-configure/**

2. Type appropriate VDP username and password in the respective fields.
3. On the **Avamar Migration** tab, review the information.

   Presence of migration history in the information indicates that VDP has been already migrated, and cannot be migrated again.

4. Click **Migrate**.

   The **Migration** wizard appears.

5. On the **Backup Jobs Migration** page:

   a. Thoroughly review the information regarding job migration and perform necessary tasks to avoid failures.

   b. Select the relevant option to enable the migration.

   c. Click **Next**.

6. On the **Destination Selection** page:

   a. Type the destination Avamar server location and the server credentials in the relevant fields.

   b. Click **Verify Authentication** to test the connection between VDP and the destination Avamar server.

   c. Click **Next**.

7. On the **Job Selection** page:

   a. From the tables on the **Backup Jobs** and **Backup Verification Jobs** tabs, select the jobs to migrate.

      All the jobs are selected by default. You can either select or deselect all the jobs by either selecting or clearing the check box on the table header.

   b. Click **Next**.

8. On the **Validation** page, review the results of validation and click **Next**.

9 On the **Ready to Complete** page, review the summary of the selected jobs and click **Migrate**.

After the migration completes, you notice the following status:

- VDP is disabled and the **Avamar Migration** tab displays the migration history.
- In the disabled state, only the following VDP features are enabled:
  - Full Image restore
  - Disk restore
  - File level restore
  - Emergency restore
- vSphere web client and VDP configuration client display the alert messages.
- Retention of the existing backups remain the same.

You can perform the following additional steps:

- To reactivate the inactivated features, click **Activate**.
- If there are application backup jobs that contain the clients that are not registered on the destination Avamar server, you must register these clients on the destination Avamar server to properly execute the newly migrated jobs. To register the clients, you must review the information in the **User Interventions** dialog box before closing the **Migration** wizard. The **User Interventions** dialog box contains a list of the clients and the domains on which you must register the clients. To check this list, click **Destination clients**.

## Troubleshooting

- After the migration completes, you can see the list of failed jobs. You can either retry or exit the migration.
- Click the **more info...** link to see additional information about the failed jobs.
- See the **avamar-migration.log** file for more information.

# Automatic Backup Verification 13

This chapter includes the following topics:


- "About Automatic Backup Verification" on page 126
- "Creating a New Backup Verification Job" on page 127
- "Editing a Backup Verification Job" on page 129
- "Cloning a Backup Verification Job" on page 129
- "Executing and Running a Backup Verification Job" on page 129
- "Monitoring Backup Verification" on page 130
- "Enabling and Disabling a Backup Verification Job" on page 130
- "Deleting a Backup Verification Job" on page 130

# About Automatic Backup Verification

Automatic Backup Verification (ABV) is a scheduled or on-demand mechanism for verification of backups that ensures the integrity of restore points. ABV has the following characteristics:

- Backups are restored to a temporary virtual machine with the following naming convention:

  **VDP_VERIFICATION_<***vm-name***> -<***unique number***>**

- Backups are restored with no network conflicts, because the NIC is always disabled during the ABV operation. Because the NIC is disabled, you cannot perform a ping test.
- After the backup verification job completes, temporary virtual machines, also referred to as validating VMs, are removed and deleted from the inventory.
- Only the last successful full image backup for any VM is verified. The Recent Tasks pane and Event log report which backup has been verified.

## Limitations

- Backup verification is not supported for the following configurations:
  - Single VMDK backups.
  - Image backups of RDM disks (physical mode). RDM virtual dependent disks are supported.
  - Application database backups of Microsoft applications (Exchange Servers, SharePoint Servers, and SQL Servers) are not supported.
  - Replicated backups.
  - Backups from imported disks.
- The verification fails if the path to the destination host is changed. If the host is moved to a new location, you must edit the verification job and select the destination host again.
- Similarly, if the datastore name is changed, you must edit the verification job to reselect the same or different datastore before you can successfully run the job again.
- In some cases, VDP fails to automatically delete the validating VM from the vCenter inventory. In this scenario, you must manually delete the validating VM.
- vSphere hosts before version 4.0 are not supported as destination hosts to where temporary VMs will be restored.

## Best Practices

### Timing and Resource Conflicts

You can take steps to avoid timing and resource conflicts when using the backup verification feature.

1. When you first install VDP, run initial full backups.
2. After the initial backups run, let the first incremental backups run.
3. Determine how long it takes for the backups to run, and schedule backup verification jobs to run after the incremental backups have completed.

### Selecting the Destination

Consider the following recommendations while selecting the destination:

- Load balance if multiple verification jobs are to be run at the same time. Limit the number of jobs to five if they are run simultaneously.
- Make sure there are sufficient resources available on the host and the datastore where the temporary VM will be restored.

### General

- Verify VMware Tools are installed on the VM at the time the VM is backed up.
- Set the heartbeat timeout interval to its optimal value, depending on the environment. Note that some VMs may take longer to send and receive the VMware Tools heartbeat than others.
- Periodically verify the availability of the destination host and a datastore. Edit the job and reconfigure the destination if needed. If the destination host or datastore is unavailable, edit the job and choose a new destination.
- Before you select Verification Script as the advanced verification option, manually run the script on the guest OS to verify it executes successfully.

## Creating a New Backup Verification Job

The backup verification job runs on demand or as part of a schedule. The Backup Verification section of the **Backup** tab allows you to create and manage backup verification jobs.

### Prerequisites

- A backup job or a restore point must exist before you create a verification job for a virtual machine. The backup job and restore type must be full image.
- VMware Tools must be installed on virtual machines at the time of backup. If no VMware Tools are found on the validating VM, the heartbeat verification will fail.
- The selected datastore must have sufficient space available.
- If you plan to use a verification script, the verification script must not be dependent on connecting to other VMs in the network.

### Procedure

1. From a web browser, access VDP.
2. Click the **Backup** tab.
3. On the **Backup** tab, click **Backup Verification**.
4. From the **Backup Verification Job Actions** menu, select **New**.

   The Create a new backup verification job wizard opens to the Virtual Machines page.
5. On the Virtual Machines page, select a virtual machine for which you want to create a verification job, and then click **Next**.
   - You can select only one virtual machine per verification. Multiple selections are not supported.
   - The virtual machine must be a part of a full image backup job or it can have restore points.
   - You can filter the virtual machines by name, if needed.
   - VMware Tools must be present on the virtual machine backups or the verification job will fail.
6. On the Verification Options page, select an option:
   - **Heartbeat Verification**: This is the default option for verification of a backup, regardless of whether you select script verification. The heartbeat verification checks whether the VMware Tools heartbeat has been received within a specific timeframe after the VM has powered on. If the VMware Tools heartbeat is received, the guest OS has booted successfully and is in a healthy state.

NOTE The **Guest OS Heartbeat** is the default option for verification.

   - **Script Verification**: This is the advanced verification option. Use script verification if you want to verify the virtual machine for the health status of applications and services that run on the guest OS. The script must be predefined and must pre-exist on the guest OS. The verification script must not be dependent on connecting to other virtual machines in the network.

If you choose to execute a script on a guest OS, supply the following information:

- **Username**: Type the user ID used to log in to the guest OS.
- **Password**: Type the password used to log in to the guest OS.
- **Confirm Password**: Retype the password.
- **Verification Script on Guest**: Type the full path to the location of the script on the guest OS.

For script configuration details, refer to "Verification Script Configuration" on page 128.

7 Click **Next**.

8 On the Destination page, select a destination:

- **Destination Path**: The destination host must be compatible with the validating virtual machine and must have sufficient resources to restore the validating virtual machine. You must select a standalone host or a host inside a cluster as a destination where backups will be temporarily restored for the purpose of verification. Resource pools and vApps are not supported as valid destinations. vSphere hosts before version 4.0 are not supported.
- **Datastore**: Depending upon the host that is selected, a list of datastores is displayed. You must select one datastore where the validating virtual machine will be restored. Make sure the selected datastore has sufficient space available.

9 Click **Next**.

10 On the Schedule page, select the schedule for the backup verification job to run. Settings made on this page determine how often and at what time of the day your verification job will run.

a **Backup verification schedule**: Specify the time intervals as daily, weekly, or monthly.

b **Start time on server**: Specify the time for the backup verification to occur on the scheduled day.

11 Click **Next**.

12 On the Job Name page, type a unique name to identify the verification job, and then click **Next**.

The verification job name can include all alphabets and numbers. The only special characters allowed are spaces, underscores, hyphens, and periods.

13 On the Ready to Complete page, review the summary of the backup verification job that you are creating. If needed, you can change the job's configuration by clicking **Back** to the appropriate page. When you are ready to save the job, click **Finish**.

**NOTE** You can also review the summary of the backup verification job from the Backup Verification section under the **Restore** tab.

14 Click **OK** when you see the message that the backup verification job was created successfully.

### Verification Script Configuration

If you plan to use a verification script, supported script formats are .bat, .cmd, .sh, and .exe. A valid script file runs by double-clicking it in the file manager or Explorer view. If the script is an unsupported format, you must enclose the execution of the script inside a supported format.

For example, a Windows Power Shell (.ps1) script cannot be called and run directly by using VDP, because the .ps format is not supported. You can, however, call the .ps script through a supported format (such as .bat), and then specify the full path to the location of the script on the guest OS. Make sure to set the execution policy to **Unrestricted** before you run the script.

The script must return a 0 or non-0 integer. If 0 is returned, the script verification succeeded. If a non-0 value is returned, the script verification failed.

## Editing a Backup Verification Job

After you create backup verification jobs, you can edit them as needed.

### Procedure

1. From a web browser, access VDP.
2. Click the **Backup** tab.
3. On the **Backup** tab, click **Backup Verification**.
4. Select the backup verification job that you want to edit, and then select **Edit** from the **Backup Verification Job Actions** menu.

   The Editing backup verification job: *job_name* wizard opens to the Virtual Machines page.
5. Step through the wizard, making changes as needed.
6. When you have completed your changes, click **Finish**.
7. Click **OK** when you see the message that changes to the backup verification job were saved successfully.

## Cloning a Backup Verification Job

You can use a verification job as a template for creating another job by highlighting the verification job and selecting **Backup Verification Job Actions** > **Clone**.

Performing the clone action launches the Cloning backup verification job wizard and uses information from the original job to automatically fill in the first three pages of the wizard (Virtual Machines, Schedule, and Retention Policy). The cloned job requires a unique name.

## Executing and Running a Backup Verification Job

After you create a backup verification job, you can invoke verification by running on-demand verification or by waiting for the schedule to start the backup verification job. The complete backup verification cycle is as follows:

- **Restore**: The latest backup for the selected VM is restored into a temporary VM that is deleted after backup verification.
- **Power On**: After the temporary VM has been restored, it will be configured to disable the NIC before it is powered on.
- **Boot OS**: Wait for the guest OS to boot completely after the VM is powered on.
- **Heartbeat Verification**: After the guest OS has booted, the appliance waits to receive the VMware Tools heartbeat from the restored VM. In an event the heartbeat is not received for any reason, the verification job fails and the backup is not in a good state.
- **Verification Script**: The script is executed only when you have selected the advanced level of verification (Verification Script). This feature runs a customized script defined and specified by the user to verify the state of applications running on the guest OS.
- **Power Off**: After the script verification completes, the VM is powered off.
- **Delete VM**: In this final step, the restored VM is deleted and the results of the verification is reported by the vCenter (in the Recent Tasks pane and in the Events log).

### Procedure

1. From a web browser, access VDP.
2. Click the **Backup** tab.
3. On the **Backup** tab, click **Backup Verification**.
4. Select the backup verification job that you want to run, and then click **Verify Now**.
5. Click **OK** when you see the message that the backup verification request or requests have been issued successfully.

## Monitoring Backup Verification

Only the last, successful backup for any VM is verified. You can use any of the methods below to check verification job results:

- vCenter Tasks/Events
- **Reports** tab. See "Viewing Information from the Reports Tab" on page 108 for more information.
- Email reports. See "Configuring the Email Notifications and Reports" on page 57 for more information.
- Client logs. You can download client logs from https://<*IP address or hostname of VDP*>:8543/vdp-configure.

## Enabling and Disabling a Backup Verification Job

After you create backup verification jobs, you can enable and disable them as needed. When you disable a backup verification job, it will not run again until you enable it.

### Procedure

1. From a web browser, access VDP. Refer to "Backup Jobs" on page 114 for instructions.
2. Click the **Backup** tab.
3. On the **Backup** tab, click **Backup Verification**.
4. Select the backup verification job that you want to enable or disable, and select **Enable/Disable** from the **Backup Verification Job Actions** menu.
5. Click **OK** when you see the message that the backup verification job has been successfully enabled or disabled.

## Deleting a Backup Verification Job

You can delete backup verification jobs when they are no longer needed.

### Procedure

1. From a web browser, access VDP. Refer to "Backup Jobs" on page 114 for instructions.
2. Click the **Backup** tab.
3. On the **Backup** tab, click **Backup Verification**.
4. Select the backup verification job that you want to delete, and select **Delete** from the **Backup Verification Job Actions** menu.
5. Click **Yes** when you are asked if you are sure you want to delete the selected job.
6. Click **OK** when you see the message that the backup verification job has been deleted successfully.

# Managing Restores 14

This chapter includes the following topic:


- "Restore Operations" on page 132
- "Selecting Backups to Restore" on page 132
- "Filtering for List of Backups" on page 133
- "Restores when Snapshots Are Present" on page 133
- "Restoring Image Backups to the Original Location" on page 133
- "Restoring Image Backups to a New Location" on page 135
- "Restoring Backups to Individual SCSI Disks" on page 136
- "Deleting a Backup from the Restore Tab" on page 137
- "Clearing all Selected Backups from the Restore Tab" on page 138

# Restore Operations

After you back up virtual machines, you can restore the backups to either the original location or an alternate location.

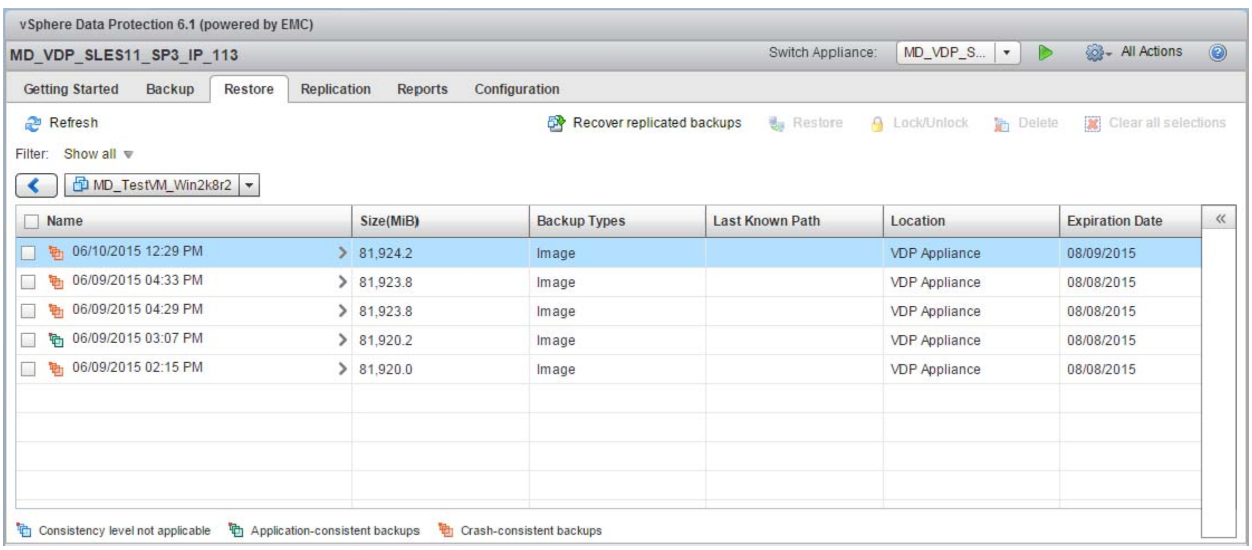
Restore operations are performed on the **Restore** tab. The **Restore** tab displays a list of virtual machines that have been backed up by the VDP appliance. By navigating through the list of backups, you can select and restore specific backups. The list shows specific icons for crash-consistent and application-consistent backups.

For icon related information, see the legend at the bottom left of the page. Before you select a backup to restore, note the crash-consistent backups and the expiration date of the backup.

**NOTE** Detection of the application-consistent backups applies only to the Windows clients. The application-consistent backups on the Linux clients appear with the **Consistency level not applicable** icon.

The information displayed on the **Restore** tab eventually becomes out-of-date. To see the most up-to-date information about the backups that are available for restore, click **Refresh**.

The following figure illustrates the **Restore** tab:




## Limitations

- The Restore wizard does not allow you to select multiple restore points for the same MSApp client. You can select only one restore point from the same client at a time.
- If the target VM has SCSI bus sharing configured, restores to that VM are not supported.

# Selecting Backups to Restore

Backups can be restored through the following options:

- Click **Restore a Backup** on the **Getting Started** tab of the VDP UI.
- From the **Restore** tab, select a restore point and click **Restore**.
- Right-click a protected virtual machine in the vCenter inventory list, and then select **All VDP Actions > Restore Rehearsal**. The Select Backup page displays a list of backups.

## Filtering for List of Backups

The list of backups that can be restored can be filtered by using drop-down arrows in the following ways:

- **Backup date**: Filtered by "is before," "is after," "is on," or "is not on"
- **Client name**: Filtered by "contains," "does not contain," "is," or "is not"
- **Location**: Filtered by location of backup

Clear the filter by clicking **Refresh** or by selecting **Show All** from the filter drop down menu.

The Select Backup page allows you to choose the virtual machines to restore.

## Restores when Snapshots Are Present

Previous versions of VDP allowed users to perform restores to the original virtual machine even if the virtual machine contained snapshots. With VDP version 5.5 and later, snapshots are not allowed on the virtual machine.

CAUTION Before you perform restores, remove any snapshots that might exist from the virtual machine. The restore job fails if the target virtual machine contains snapshots.

Refer to the following Knowledge Base articles about how to properly use snapshots:

- http://kb.vmware.com/kb/1025279
- https://community.emc.com/thread/145249?start=0&start=0

## Restoring Image Backups to the Original Location

You can restore backups manually by using the Restore backup wizard.

There are three scenarios where, if the individual disks are selected to be restored instead of the entire VM, you cannot restore to the original location:

- The original disk is marked as independent-persistent.
- The original disk has been removed from the target VM.
- The original disk has been deleted from the target VM.

NOTE A restore job of the same disk or disks from two different timestamps is not permitted. If you attempt to restore a disk that has been backed up with two different timestamps, you are presented with the option of removing the duplicated hard disk. The restore will not proceed until the duplicated hard disk is removed.

### Prerequisites

- VDP is installed and configured on your vCenter server.
- You are logged in to the vSphere web client and connected to the VDP appliance.

### Procedure

1. From a web browser, access VDP.
2. Click the **Restore** tab.
3. If necessary, filter the backups to narrow your search.
4. Select a virtual machine listed in the Name column. When you click on a virtual machine, it expands to list the backups that have been performed. You can select one or more backups, or you can click a backup to locate the disk to restore.

NOTE The client (virtual machine) name is renamed to append a string of random characters in the **VDP_IMPORT** domain on the **Restore** tab if storage is imported from a different VDP appliance during initial configuration.

5 Select the checkbox beside one or more items to select them for restore.

6 Click **Restore** to start the Restore backup wizard.

7 On the **Select Backup** page:

   a Review the list of selected backups for crash consistent backups. If you do not want to restore the crash consistent backups, remove them.

   b Click **Next**.

8 On the **Set Restore Options** page:

   a Leave the **Restore to original location** option selected. If the vmdk file still exists at the original location, it is overwritten.

**NOTE** If the virtual disk on the original VM has been removed or deleted, the restore to original location option is not allowed. The VMDK must be restored to a new location.

   b If you want to restore the virtual machine along with its configuration, select **Restore virtual machine along with configuration**.

   c Click **Next**.

9 On the **Ready to complete** page:

   a Review the summary of your restore request.

   The summary displays the following information:

   - The number of machines that will be replaced or restored to their original location, and created or restored to a new location
   - A checkbox for your consent to proceed with the restore if the list of selected backups to restore contains crash consistent backups

     If you select the checkbox to proceed with the crash consistent restore, the restore sometimes fails.

   If you want to change any of the restore settings, click either the **Back** button on the wizard page or click the appropriate numbered step title on the left side of the wizard page to go to the appropriate page.

   b Click **Finish** to start the restore operation.

   A message that states that the restore has been successfully initiated appears.

   c Click **OK**.

10 Monitor the progress of the restore in the **Recent Tasks** panel.

**NOTE** If you selected **Reconnect NIC** during the restore process, confirm the network configuration for the newly-created virtual machine. It is possible that the new virtual machine NIC is using the same IP address as the original virtual machine, which will cause conflicts.

# Restoring Image Backups to a New Location

You can manually restore backups by using the Restore a backup wizard. On the **Set Restore Options** page of the Restore a backup wizard, you can specify the destination to restore the image backup.

**NOTE** If you have included physical Raw Device Mapping (RDM) disks in an image backup, you cannot restore the backup to a new virtual machine. If you back up a virtual machine that has both virtual disks and physical RDM disks, the backup successfully processes the virtual disks only. You can restore the backed up data to only either the source virtual machine or redirect it to another (already existing) virtual machine. Because the backup does not process the physical RDM disks, you cannot restore the data that is present on the physical RDM disks. So, you cannot restore an image backup to a new virtual machine. https://www.emc.com/collateral/TechnicalDocument/docu59591.pdf provides more information.

### Prerequisites

- VDP is installed and configured on your vCenter server.
- You are logged in to the vSphere web client and connected to the VDP appliance.

### Procedure

1. From a web browser, access VDP.
2. Click the **Restore** tab.
3. If necessary, filter the backups to narrow your search.
4. Select a virtual machine listed in the Name column. When you click on a client (virtual machine), it expands to list the backups that have been performed. You can select one or more backups, or you can click a backup to drill down further until you locate the disk or application that you want to restore.

**NOTE** The client name that appears in the Name column is renamed to append a string of random characters if the storage has been imported from a different VDP appliance during initial configuration.

5. Select the checkbox beside one or more items to select them for restore.
6. Click **Restore** to start the Restore backup wizard.
7. On the **Select Backup** page, click the backup that you want to restore and click **Next**.
8. On the **Set Restore Options** page:
   a. Clear the **Restore to Original Location** box to set the restore options for each backup that you are restoring to a new location.
   b. Specify the following fields:
      - **New VM Name**: Type a new name for the restored VM.
      - **Destination**: Click **Choose** and select the new destination.
      - **Datastore**: Select the datastore in which the VM will be restored.
   c. Click **Next**.
9. On the **Ready to Complete** page:
   a. Review the summary of your restore requests.

      This summary identifies how many machines will be replaced (or restored to their original location) and how many will be created (or restored to a new location).

      If you want to change any of the settings for your restore request, click the **Back** button to return to the appropriate screen, or click the appropriate numbered step title on the left side of the wizard screen.
   b. Click **Finish** to start the restore operation.

      A message that states that the restore has been successfully initiated appears.

c Click **OK**.

10 Monitor the progress of the restore in the **Recent Tasks** panel.

**NOTE** If you selected **Reconnect NIC** during the restore process, confirm the network configuration for the newly-created virtual machine. It is possible that the new virtual machine NIC is using the same IP address as the original virtual machine, which will cause conflicts.

## Disabling the vMotion Feature Before Performing Restores to a New Location

vSphere vMotion is a feature that enables live migration of the running virtual machines from one physical server to another. Ensure that the vMotion operations do not run on the VDP appliance while you perform backups or restores. "Storage vMotion operations are not allowed during active backup operations" on page 201 provides information.

Disable the vMotion feature at the VM and datastore levels, in the backup window and before you perform a restore to a new location.

1 From a web browser, access the vSphere web client:

 **https://**<IP_address_vCenter_Server>**:9443/vsphere-client/**

2 Select **vCenter** > **Hosts and Clusters**.

3 On the **DRS** tab, select **Manual** for all the DRS clusters that are listed in the left menu.

4 Select **Inventory** > **Datastores**.

5 For each datastore on which VMs, proxies, and VDPs reside:

 a Select the datastore.

 b Select **Edit Settings**.

 c On the left menu, under **General**, click **SDRS Automation**.

 d In the right panel, select **No automation (Manual)**.

 e Select **Virtual Machine Settings**.

 f Ensure that all VMs show either **Disabled** or **Manual** for the SDRS Automation.

# Restoring Backups to Individual SCSI Disks

You can restore backups to individual SCSI disks by using the Restore backup wizard. On the Set Restore Options page of the Restore a backup wizard, you can specify to where you want the individual SCSI disks restored.

**NOTE** SCSI ID does not support multiple restore requests of different VMs. Although you can initiate multiple restore operations, only the first restore is successful. When restoring multiple disks as new disks on either the original or existing VM, all the disks to be restored must come from the same backup and the same VM.

### Prerequisites

- VDP is installed and configured on your vCenter server.
- You are logged in to the vSphere web client and connected to the VDP appliance.

### Procedure

1 From a web browser, access VDP.

2 Click the **Restore** tab.

3 If necessary, filter the backups to narrow your search.

4 Select a virtual machine listed in the Name column. When you click on a client (virtual machine), it expands to list the backups that have been performed. You can select one or more backups, or you can click a backup to drill down further until you locate the disk or application that you want to restore.

**NOTE** The client name that appears in the Name column is renamed to append a string of random characters if the storage has been imported from a different VDP appliance during initial configuration.

5 Select the checkbox beside one or more items to select them for restore.

6 Click **Restore** to start the Restore backup wizard.

7 On the **Select Backup** page, click the backup to restore and click **Next**.

8 On the **Set Restore Options** page, clear **Restore to original location** to set the restore options for each backup that you are restoring to a new location.

a Specify the following information:

- **Destination**: To select a new destination, click **Choose.** You can select a new location destination container (vApp, Resource pool, host, or datacenter) where the backup will be restored, or you can accept the default destination, which is the original location of the existing VM.
- **New VM Name**: The **New VM Name** field auto-populates with the existing VM name. You can modify this field to give the VM a new name if you are restoring to a container. If you are restoring to an existing VM, you cannot modify the name.
- **Datastore**: Lists the datastore in which the first disk currently resides. If you are restoring a disk to an existing VM, this field is not editable. If you are restoring a disk to a new container, select the datastore in which the VM will be restored.
- **Disk ID**: Lists the SCSI disk ID slots available as restore targets. The list shows only empty SCSI slots of the SCSI controllers that are currently attached to the virtual machine. Select a SCSI virtual disk slot as the restore target from the list.

  You cannot restore to IDE-configured virtual disks. Only SCSI virtual disks are supported. A maximum of 15 disks are allowed on a SCSI controller. Slot 7 is reserved and unavailable.

**NOTE** The Restore backup wizard does not automatically add a new controller if there are not enough SCSI slots available. You must add a new SCSI controller before initiating the disk restore.

b Click **Next**.

9 On the **Ready to Complete** page:

a Review the summary of your restore requests.

This summary identifies how many machines will be replaced (or restored to their original location) and how many will be created (or restored to a new location).

If you want to change any of the settings for your restore request, click the **Back** button to return to the appropriate screen, or click the appropriate numbered step title on the left side of the wizard screen.

b Click **Finish** to start the restore operation.

A message that states that the resore has been successfully initiated appears.

c Click **OK**.

10 Monitor the progress of the restore in the **Recent Tasks** panel.

## Deleting a Backup from the Restore Tab

VDP deletes backups according to the retention policies that were set in the backup jobs. However, you can manually delete backups from the **Restore** tab by selecting the backup jobs for deletion and clicking the **Delete** icon.

## Clearing all Selected Backups from the Restore Tab

1. From the **Manual Restore** tab, select the backups you want to clear from the list of backups, and then click **Clear All Selections**.
2. Click the **Refresh** button to update the data in the **Restore** tab.

# 15 Replication

This chapter includes the following topics:


- "Replication Jobs" on page 140
- "Creating a Replication Job" on page 143
- "Managing Destinations" on page 147
- "Editing a Replication Job" on page 147
- "Cloning a Replication Job" on page 147
- "Deleting a Replication Job" on page 148
- "Enabling or Disabling a Replication Job" on page 148
- "Viewing Status and Replication Job Details" on page 148
- "Running Existing Replication Jobs Immediately" on page 148
- "Replication Back to the Source" on page 148
- "Replication Recovery Compatibility" on page 149
- "Enabling or Disabling Replication Recovery" on page 150
- "Replication Recovery" on page 150
- "Multi-Tenancy" on page 151

# Replication Jobs

Replication enables you to avoid data loss if the source VDP appliance fails because copies of the backups are available on the destination target. Replication jobs determine which backups are replicated, and when and to where the backups are replicated. With scheduled or ad hoc replication jobs for clients that have no restore points, only the client gets replicated on the destination server. Backups created with the VDP appliance can be replicated to another VDP appliance, to an Avamar server, or to a Data Domain system.

## Replication Compatibility

Table 15-20 and Table 15-21 indicate which backups can and cannot be replicated, depending on which VDP product was used to create them.

The following abbreviations are used in these tables:

- **VDP-A**: VDP Advanced appliance
- **RTI**: Replication Target Identity

IMPORTANT VDP Advanced applies to VDP 5.8 and earlier. RTI applies only to VDP 5.8.

- **DD**: Data Domain System
- **N**: No, the target is not supported
- **Y**: Yes, the target is supported
- **(R)**: Recommended
- **(NR)**: Not Recommended
- **(HP)**: Hashed Password

**Table 15-20.** Replication Source Matrix - Part 1

| Backups created with this product... | Can be replicated to these targets... | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | VDP 5.1.x | VDP 5.5.1.x | VDP 5.5.5.x | VDP-A 5.5.5.x | VDP-A 5.5.5.x + DD | VDP 5.5.6.x | VDP-A 5.5.6.x | VDP-A 5.5.6.x + DD | VDP 5.8.x | VDP-A 5.8.x | VDP-A 5.8.x + DD | VDP 6.0.x | VDP 6.0.x + DD | VDP 6.1.x | VDP 6.1.x + DD |
| **VDP 5.1.x** | N | N | N | N | N | N | N | N | N | N | N | N | N | N | N |
| **VDP 5.5.1.x** | N | N | N | N | N | N | N | N | N | N | N | N | N | N | N |
| **VDP 5.5.5.x** | N | N | N | N | N | N | N | N | N | N | N | N | N | N | N |
| **VDP-A 5.5.5.x** | N | N | N | Y (R) | Y (R) | N | Y | Y | N | Y (HP) | Y (HP) | Y (HP) | Y (HP) | Y (HP) | Y (HP) |
| **VDP-A 5.5.5.x + DD** | N | N | N | N | Y (R) | N | N | Y | N | N | Y (HP) | N | Y (HP) | N | Y (HP) |
| **VDP 5.5.6.x** | N | N | N | N | N | N | N | N | N | N | N | N | N | N | N |
| **VDP-A 5.5.6.x** | N | N | N | Y | Y | N | Y | Y | N | Y (HP) | Y (HP) | Y (HP) | Y (HP) | Y (HP) | Y (HP) |
| **VDP-A 5.5.6.x + DD** | N | N | N | N | Y | N | N | Y | N | N | Y (HP) | N | Y (HP) | N | Y (HP) |
| **VDP 5.8.x** | N | N | N | N | N | N | N | N | N | N | N | N | N | N | N |
| **VDP-A 5.8.x** | N | N | N | Y | Y | N | Y | Y | N | Y (R) | Y | Y (R) | Y | Y (R) | Y |
| **VDP-A 5.8.x + DD** | N | N | N | N | Y | N | N | Y | N | N | Y (R) | N | Y (R) | N | Y (R) |
| **RTI 5.8.x** | N | N | N | Y | Y | N | Y | Y | N | Y (R) | Y | Y (R) | Y | Y (R) | Y |

**Table 15-20.** Replication Source Matrix - Part 1 (Continued)

| Backups created with this product... | Can be replicated to these targets... | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | VDP 5.1.x | VDP 5.5.1.x | VDP 5.5.5.x | VDP-A 5.5.5.x | VDP-A 5.5.5.x + DD | VDP 5.5.6.x | VDP-A 5.5.6.x | VDP-A 5.5.6.x + DD | VDP 5.8.x | VDP-A 5.8.x | VDP-A 5.8.x + DD | VDP 6.0.x | VDP 6.0.x + DD | VDP 6.1.x | VDP 6.1.x + DD |
| **RTI 5.8.x + DD** | N | N | N | N | Y | N | N | Y | N | N | Y (R) | N | Y (R) | N | Y (R) |
| **VDP 6.0.x** | N | N | N | Y | Y (R) | N | Y | Y (R) | N | Y | Y (R) | Y | Y (R) | Y | Y (R) |
| **VDP 6.0.x + DD** | N | N | N | N | Y | N | N | Y | N | N | Y | N | Y | N | Y |
| **VDP 6.1.x** | N | N | N | Y | Y (R) | N | Y | Y (R) | N | Y | Y (R) | Y | Y (R) | Y | Y (R) |
| **VDP 6.1.x + DD** | N | N | N | N | Y | N | N | Y | N | N | Y | N | Y | N | Y |

**Table 15-21.** Replication Source Matrix - Part 2

| Backups created with this product... | Can be replicated to these targets... | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | RTI 5.8.x | RTI 5.8.x + DD | VDP 6.0.x | VDP 6.0.x + DD | Avamar SP1 6.1.1.87 | Avamar SP2 6.1.2.47 | Avamar 7.0.0.427 | Avamar 7.0.1.56 | Avamar 7.1.x.x | Avamar 7.1.x.x + DD |
| **VDP 5.1.x** | N | N | N | N | N | N | N | N | N | N |
| **VDP 5.5.1.x** | N | N | N | N | Y | Y | Y | Y | Y | Y |
| **VDP 5.5.5.x** | N | N | N | N | Y | Y | Y | Y (R) | Y | Y |
| **VDP-A 5.5.5.x** | Y (HP) | Y | Y | N | Y | Y | Y (R) | Y (R) | Y | Y |
| **VDP-A 5.5.5.x + DD** | N | Y (HP) | N | Y (HP) | Y (NR) | Y (NR) | Y (NR) | Y (NR) | Y | Y |
| **VDP 5.5.6.x** | N | N | N | N | Y | Y | Y (R) | Y | Y | Y |
| **VDP-A 5.5.6.x** | Y (HP) | Y (HP) | Y (HP) | N | Y | Y | Y (R) | Y (R) | Y | Y |
| **VDP-A 5.5.6.x + DD** | N | Y (HP) | N | Y (HP) | Y (NR) | Y (NR) | Y (NR) | Y (NR) | Y | Y |
| **VDP 5.8.x** | N | N | N | N | Y | Y | Y | Y | Y (R) | Y |
| **VDP-A 5.8.x** | Y (R) | Y | Y (HP) | N | Y | Y | Y | Y | Y (R) | Y |
| **VDP-A 5.8.x + DD** | N | Y (R) | N | Y (HP) | Y (NR) | Y (NR) | Y (NR) | Y | Y | Y |
| **RTI 5.8.x** | Y (R) | Y | Y (HP) | N | Y | Y | Y | Y | Y (R) | Y |
| **RTI 5.8.x + DD** | N | Y (R) | N | Y (HP) | Y (NR) | Y (NR) | Y (NR) | Y (NR) | Y (NR) | Y (R) |
| **VDP 6.0.x** | Y | Y (R) | Y (HP) | Y (R) | Y (HP) | Y (HP) | Y (HP) | Y (HP) | Y (HP) | N |
| **VDP 6.0.x + DD** | N | Y | N | Y (HP) | N | N | N | N | N | Y (HP) |

**Table 15-21.** Replication Source Matrix - Part 2

| Backups created with this product... | Can be replicated to these targets... | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | **RTI 5.8.x** | **RTI 5.8.x + DD** | **VDP 6.0.x** | **VDP 6.0.x + DD** | **Avamar SP1 6.1.1.87** | **Avamar SP2 6.1.2.47** | **Avamar 7.0.0.427** | **Avamar 7.0.1.56** | **Avamar 7.1.x.x** | **Avamar 7.1.x.x + DD** |
| **VDP 6.1.x** | Y | Y (R) | Y (HP) | Y (R) | Y (HP) | Y (HP) | Y (HP) | Y (HP) | Y (HP) | N |
| **VDP 6.1.x + DD** | N | Y | N | Y (HP) | N | N | N | N | N | Y (HP) |

## Replication and Data Domain

If the source VDP appliance has a Data Domain system as its backup target, the replication destination VDP appliance must also have a Data Domain system. Similarly, if replicating from VDP to an Avamar server, the Avamar server must have a Data Domain system.

**NOTE** Data Domain Boost 2.6 and Data Domain 5.4.1.1 or later, 5.5.x, and 5.6 are supported.

The Replication job fails if Data Domain and Avamar backup targets are combined into a single replication job. You must either configure all Data Domain clients or all Avamar clients as the backup targets.

## Best Practices for Replication

- Because only completed client backups are replicated, make every effort to schedule replication during periods of low backup activity. This ensures that the greatest number of client backups replicate during each replication session.
- If you change the user ID or password for the root account on the replication target server, you must update the destination user ID and password on the source server with the new password.
- Using replication destination management, you can update the information for one or more replication jobs that are associated with the same replication destination server.
- Replication of dynamic or non-static data is not supported. Therefore, you must run the replication during a period of low backup activity.
- You cannot run simultaneous replication and/or recovery of multiple clients if simultaneous, multiple backup and/or restore operations are running.
- Though replication between servers of different versions is supported, ensure that the VDP version on the destination server is either the same or later than the VDP version on the source server.
- While creating the replication job policy, optimize the size of the replication group so that all the clients successfully replicate during each scheduled replication operation. If the group size is very large, replication takes more time to complete. So, split the group into two smaller groups that you can independently schedule.
- Multiple replication jobs, which are part of a single job, for different VMs run do not parallely run. To change this behavior, add the **com.vmware.vdp.option.replicate.maxtreams** property to the **/usr/local/vdr/etc/vdp-options.properties** file.

  For example, com.vmware.vdp.option.replicate.maxstreams=8

## Limitations

- When a replication job starts, it can process only quiescent, static data that resides on the source server. Therefore, any operation that writes data to the source server and has not fully completed (for example, an in-process backup job) will not be part of that replication job. The data will, however, be replicated during the next replication operation.
- On the source VDP appliance, as the number of backups for replicated clients increases, the time required to browse each client increases.
- You can replicate the imported backups only by selecting the replicated backup types.

  For imported backups, the import process moves the clients and accounts from their current account path to /REPLICATE/VDP_IMPORTS/. The name of the domain at this level includes a timestamp, which is appended to the end of the name. The import process also deletes all jobs. To replicate these backups, you must select replicated backups for the type, and then select imported backups.

## Defining Backup Types for a Replication Job

You specify the retention policy and backup schedule when you create backups. Consider these factors when defining the backup types you will use for the replication job.

- Retention policy. Refer to "Setting the Retention Policy" on page 115 for detailed information.
- On-demand or scheduled backup types:
  - If any backup job is run by using the **Backup Now** option, it is considered to be an on-demand backup and is not associated with any backup type. To replicate this backup, you must select the **User Initiated** backup type (Step 6 of the Create a new replication job wizard).
  - To schedule the backup, specify the scheduling options (Step 10 of the Create a new replication job wizard).

# Creating a Replication Job

You create replication jobs by using the Create a new replication job wizard.

**NOTE** Clients or restore points that have already been replicated from a different source server are available in the Create a new replication job wizard.

### Prerequisites

- VDP is installed and configured on your vCenter server.
- You are logged in to the vSphere web client and connected to the VDP appliance.

## Procedure

1. From a web browser, access VDP. Refer to "Accessing VDP" on page 106 for instructions.
2. Click the **Replication** tab.

   The **Replication** tab displays a list of the replication jobs that have been created. The following table lists the columns and their description:

   **Table 15-22.** Column descriptions for the **Replication** tab

| Column | Description |
|---|---|
| Name | The name of the replication job. |
| State | The state of the replication job. |
| Destination | The location where the client backups are replicated. |
| Last Run Time | The last time the replication job ran. |
| Duration | How long the replication took to complete the last time the job ran. |
| Next Run Time | When the replication job is scheduled to run again. |
| # of Clients | The number of clients whose backups are being protected and replicated in the job. This value changes only when the user adds or removes clients from a replication job using the Edit feature. |

3. From the **Replication Job Actions** menu, select **New** to start the Create a new replication job wizard.
4. On the Select Type page, choose whether to replicate guest images (local backups) or replicated backups, and then click **Next**.

   The Select Type page appears. On this page, the appropriate clients appear, based on the type of replication you selected.

   If the type is replicated backups:

   - The VDP appliance displays replicated backups and recovered backup clients as options to replicate to another target server.
   - The client table shows the source path in the Source Path column, rather than the last known path.

**NOTE** Both guest images and application backup options are available.

5. On the Select Clients page, perform one of the following tasks:

   - To replicate all client backups, click **All clients**, and then click **Next**.
   - To replicate backups from specific clients only, click **Select clients individually**, and then select the type from the **Type** list. Options include Image, MS SQL Servers, MS Exchange Servers, and MS SharePoint Servers.

   **NOTE** Both regular and retired VM backups are supported for replication. If a retired VM is re-added as a regular VM, the system lists the VM twice with an identical name. The retired VM name is appended with a suffix. When selecting clients, select the regular VM without the suffix.

   If you choose **Select clients individually**, you can select one or more clients. If preferred, you can filter the clients before you make any selections. To filter clients:

   a. Beside **Filter**, click **Show All** and select **Client**.

      To filter by client name, select **Name**. The following information appears for the vCenter Client.

      - **Name:** "Is," "Is not," "Contains," or "Does not contain" filters used to query the client name.
      - **Status:** Values are Powered On, Powered Off, Suspended, Activated, or Not Activated.
      - **Client Type**: The type of client.

b Click **Next**.

The **Backup Selection** page appears. On this page, you can limit the number of backups that are replicated when the job runs. If you choose not to select backup options, every backup for the selected virtual machines will be replicated.

6 On the **Backup Selection** page of the **Create a new replication job** wizard:

a Select a **Backup Type**:

**Daily:** Only daily backups will be replicated.

**Weekly:** Only weekly backups will be replicated once.

**Monthly:** Only monthly backups will be replicated.

**Yearly:** Only yearly backups will be replicated.

**User Initiated:** Only those backups that were user-initiated will be replicated.

**NOTE** User-initiated backups do not retain advanced retention options. These backups must be flagged as a separate backup type.

b Specify the **Maximum backups to replicate per client**:

**No Limit:** When this option is selected, all existing backups for a client that meet the **Backup Type** criteria will be replicated. The number of backups is unlimited.

**Number of Backups:** When this option is selected, the backups to be replicated are based on chronological order. The most recent backup is selected, regardless of whether it is an on-demand backup or a scheduled backup. The maximum number of backups is 999.

c Specify the **Date Restrictions**:

**None:** All backups that meet the **Backup Type** and **Maximum backups to replicate per client** criteria will be replicated. There are no other restrictions.

**Last:** Select a number and a time unit. This option restricts the selection of backups by including only those backups that were created during the specified number of days, weeks, months, or years.

**By Range:** Select a **From** date and time, and select a **To** date and time. You can specify from a specific date forward, until a specific date, or between two dates.

d Click **Next**.

The **Destination** page appears. On this page, you specify connection information for the destination where the client backups will be replicated.

You can use an Avamar server as a replication destination. To do this, supply the Avamar server's IP address, port, and login credentials on the Destination page.

**NOTE** If you change the name of the VM client, VDP displays the renamed client in the Name column on the Create a new replication job wizard. If you perform a replication of a renamed VM client and the destination is an Avamar server or Avamar Virtual Edition (AVE), the changed name is not reflected in Avamar. The Avamar server displays the older name that was previously registered before the name was changed. This is a known issue.

7 On the **Destination** page of the Create a new replication job wizard, provide the following information:

**NOTE** All references to **Destination** refer to either the Avamar server or the VDP appliance to which the backup data is being replicated.

- **Hostname or IP**: The hostname or IP address of the destination.
- **Port**: The port number over which VDP communicates with the destination. The default value is 29000, which is the standard port for SSL encrypted replication.
- **Username**: The username used to log in to the destination.
- **Password**: The password used to log in to the destination.

- **Path**: The unique name that identifies the domain, used for multi-tenancy configurations. Refer to "Multi-Tenancy" on page 151 for details about multi-tenancy configuration.

For replication and recover replicated backup operations, use the **repluser** username. The credentials of the **repluser** username are kept in sync with the root user.

8. Click **Verify Authentication** to test the connection between VDP and the destination.
9. Click **Next**.

   The **Schedule** page appears. On this page, you specify how often backups will be replicated and what time of the day the replications will occur.

10. On the **Schedule** page of the Create a new replication job wizard:

    a. Select one of the schedule options:
       - **Daily**: Select this option to replicate the backups every day.
       - **Weekly performed every**: Select this option, and select a day to replicate the backups on that day every week.
       - **The .... of every month:** Select this option, and select a number and a day to replicate the backups on that day of every month.

    b. Select a **Start Time on Server** to specify the time that the replication will take place on the scheduled day.

       Best practice: Because only completed client backups are replicated, you should make every effort to schedule replication during periods of low backup activity. This ensures that the greatest number of client backups replicate during each replication session.

    c. Click **Next**.

       The Retention page appears. On this page, you specify when replicated backups will expire on the destination machine.

11. On the **Retention** page of the Create a new replication job wizard:

    a. To use each backup's current expiration date, select **Keep the current expiration for each backup**.

    b. To specify expiration dates based on backup type, select **Set expiration by backup type**, and select the number of days, weeks, months, or years for each type.

    c. To keep the replication job forever, select **Keep forever**.

    d. Click **Next**.

       The Name page appears. On this page, you name the replication job.

12. On the **Name** page of the Create a new replication job wizard:

    a. Type a name for the replication job.

       The replication job name must be unique and can be up to 255 characters long. The following characters cannot be used in the job name: ~!@$^%(){}[]|,`;#\/:*?<>'"&. In addition, diacritical characters cannot be used (for example: â, é, ì, ü, and ñ).

    b. Click **Next**.

       The Ready to complete page appears. On this page, you can review a summary of the replication job that you are creating before saving the job.

13. On the **Ready to complete** page of the Create a new replication job wizard:

    a. Review the information.

    b. Click **Finish** to create the job.

# Managing Destinations

You can select existing replication jobs from the **Replication** tab and then change the destination connection information for all of them using the 3-step Manage Destination wizard.

### Best Practice

All replication jobs associated with the same specific replication destination server should be updated, rather than selecting a combination of replication jobs with various destination servers.

### Procedure

1. From a web browser, access VDP. Refer to "Accessing VDP" on page 106 for instructions.
2. Click the **Replication** tab.

   The **Replication** tab displays a list of the replication jobs that have been created
3. Highlight the replication job, and select **Replication job actions > Manage Destination**.

   The Manage Destination Wizard appears.
4. On the Replication Jobs page, click the replication job to update its associated destination, and then click **Next**. You can select multiple jobs.
5. On the **Destination** page of the Manage Destination wizard, provide the following information:

NOTE All references to **Destination** refer to either the Avamar server or to the VDP appliance to which the backup data is being replicated.

- **Hostname or IP**: The hostname or IP address of the destination.
- **Port**: The port number over which VDP communicates with the destination. The only allowable port is 29000. This port is the standard port for SSL encrypted replication.
- **Username**: The username used to log in to the destination.
- **Password**: The password used to log in to the destination.
- **Path**: The unique name that identifies the domain. This field is used for multi-tenancy configurations. Refer to "Multi-Tenancy" on page 151 for details about multi-tenancy configuration.

For Replication and Recover Replicated backup operations, use the **repluser** username. The credentials of the **repluser** username are kept in sync with the root user.

6. Click **Verify Authentication** to test the connection between VDP and the destination.
7. Click **Next**.
8. On the Ready to Complete page, review the destination that will be assigned to the selected replication jobs. Click **Finish** to update the replication job, or click **Back** to make changes.

# Editing a Replication Job

You can edit a replication job by highlighting it and selecting **Replication job actions** > **Edit**.

# Cloning a Replication Job

You can use a replication job as a template for creating a different job. Highlight the replication job, and select **Replication job actions** > **Clone**.

Performing the clone action launches the Cloning replication job wizard and uses information from the original job to automatically fill in the information. The cloned job requires a unique name. You can modify any of the settings that were copied from the original job.

## Deleting a Replication Job

You can delete a replication job by highlighting it and selecting **Replication job actions** > **Delete**.

NOTE You can select multiple replication jobs to delete. Deleting a replication job will decrement the count of replication jobs associated with a particular replication destination. If the replication jobs being deleted result in replication destinations with no associated replication jobs, then you are given an option to remove replication destinations as part of the delete request.

## Enabling or Disabling a Replication Job

If you want to temporarily stop a replication job from running in the future, you can disable it. You can edit and delete disabled replication jobs. VDP will not run a disabled job until it has been enabled.

You can enable or disable a replication job by highlighting the job and selecting **Replication job actions** > **Enable/Disable**.

## Viewing Status and Replication Job Details

The **Replication** tab displays a list of replication jobs that have been created with VDP. You can see the details of a replication job by clicking the job. The details are displayed in the Replication Job Details pane:

- **Name:** The name of the replication job.
- **State:** The state of the replication job.
- **Destination:** Where the backups specified in the job were replicated.
- **Clients:** A list of the clients whose backups are replicated by the job.
- **Last Run time**: The last time the replication job ran.
- **Duration**: How long the replication took to complete the last time the job ran.
- **Next Run time**: The date and time the job is next scheduled to run.

## Running Existing Replication Jobs Immediately

You can run a replication job immediately by highlighting the job and clicking **Replicate Now**.

## Replication Back to the Source

A replication job can be set up to replicate backup data from one VDP appliance to another VDP appliance.

### Node Structure for Recovered Backups

The first time a valid recovery action initiates using the source server, the recovery action creates the source server node under the /REPLICATE link. All recovered backups are displayed under that /REPLICATE link on the **Restore** tab of the source VDP appliance.

After a successful replication from the source VDP appliance to a replication target server, when a recovery occurs from the replication source server for replicated backups located at a target server, the fully qualified domain name (FQDN) of the source server appears under the /REPLICATE link on the **Restore** tab. Recovered backups do not display the target server FQDN from which it recovered.

### Node Structure of Backups Replicated Again

When a user replicates a replicated backup, the replication source node is not displayed under the /REPLICATE link on the replication target server.

After a successful replication of replicated backups, the **Restore** tab of the subsequent target server displays the parent source server information (on the source server where the virtual machine was originally backed up) under the /REPLICATE link on the **Restore** tab.

After a successful replication of replicated backups from server B to server C (where the virtual machine was originally backed up using server A), the **Restore** tab on the VDP appliance of server C displays the information for server A under the /REPLICATE link, rather than displaying information for the actual replication source server B.

## Replication Destinations

The following are examples of where the replication destination is required:

- A backup that exists on the local VDP appliance is replicated to a remote replication target, and then the backup is removed or deleted from the local appliance. With the **Replication Back to the Source** feature, you can browse the backups that reside on a remote destination and recover specific backups by copying them back to the local appliance. After the backup is recovered back on the local appliance, it can be restored with the usual process.
- A new VDP appliance must be installed to replace a corrupted VDP. Specific backups that have already been replicated to a replication destination must be recovered. The **Replication Back to the Source** feature allows you to connect to a replication destination, browse the backups that reside there, and recover specific backups by copying them back to the local appliance. After the backup is recovered back on the new VDP appliance, it can be restored with the usual process.

NOTE The Recover wizard is not available on the Avamar server. The ability to replicate back to the source is not an option on the Avamar server.

# Replication Recovery Compatibility

Table 15-23 lists the supported replication recovery target and replication recovery source servers.

**Table 15-23.** Replication Recovery Compatibility Matrix

| Replication Recovery Target (Recover From) | Replication Recovery Source (Recover To) | | | | |
|---|---|---|---|---|---|
| | VDP 5.8 | VDP Advanced 5.8 | VDP Advanced 5.8 Target Identity | VDP 6.0 | VDP 6.1 |
| VDP 5.8.0.x | No | No | No | Yes | Yes |
| VDP Advanced 5.8.0.x | No | Yes | Yes | Yes | Yes |
| VDP Replication Target Identity 5.8.0.x | No | Yes | Yes | Yes | Yes |
| VDP 5.5.1.356 | No | No | No | No | No |
| VDP 5.5.5.180 | No | No | No | No | No |
| VDP Advanced 5.5.5.180 | No | Yes | Yes | Yes | Yes |
| VDP 5.5.6.56 | No | No | No | No | No |
| VDP Advanced 5.5.6.56 | No | Yes | Yes | Yes | Yes |
| Avamar server 7.0.x | Yes | Yes | Yes | Yes | Yes |
| Avamar Virtual Edition (AVE) 6.0.x | Yes | Yes | Yes | Yes | Yes |
| Avamar server / AVE 7.1.x | Yes | Yes | Yes | Yes | Yes |
| AVE 7.2.x | Yes | Yes | Yes | Yes | Yes |
| VDP 6.0 | No | Yes | Yes | Yes | Yes |
| VDP 6.1 | No | Yes | Yes | Yes | Yes |

IMPORTANT VDP Advanced applies to versions 5.8 and earlier. Replication Target Identity applies only to VDP 5.8.

# Enabling or Disabling Replication Recovery

For VDP versions 5.5 and later, port 29000 is used when replicating from a VDP appliance to a replication target destination, such as an Avamar storage server or a Data Domain system.

To enable the replication recovery feature, port 29000 must be open on the replication source server. Opening port 29000 allows replicated backups to the replication target destination. If port 29000 is closed, the replication recovery becomes disabled and you are unable to verify the destination as a valid replication recovery target.

By default port 29000 is open on a VDP appliance. Refer to standard Linux server documentation for procedures on how to open and close port 29000.

# Replication Recovery

On a newly installed VDP, you can specify a destination (a machine to where backups have been replicated) by using the Create a new replication job wizard. After the destination is added, you can recover and restore any replicated backups.

**NOTE** With the current implementation of replication recovery, if a virtual machine client is renamed on a local VDP appliance, then that name change will not be propagated to backups on a remote VDP Advanced appliance that is used as a replication target. The new virtual machine name is reflected in the local restore points, but it is not reflected in the destination.

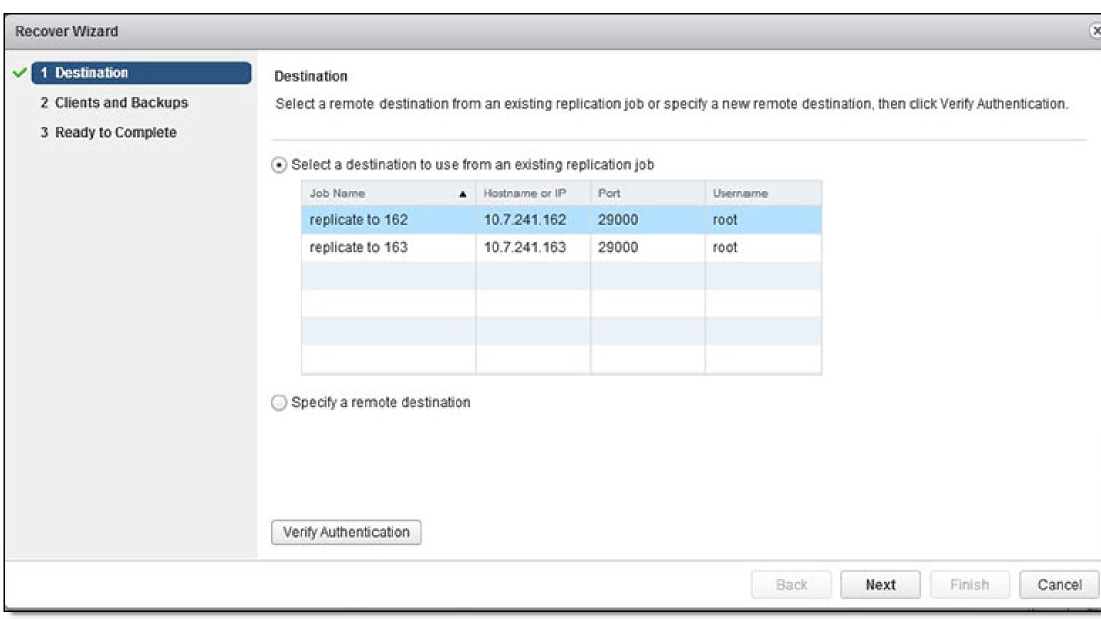
The **Select a replication target** option is disabled after a vCenter server change and is not enabled until a new replication job is created after the vCenter server change.

1 From a web browser, access VDP. Refer to "Accessing VDP" on page 106 for instructions.

2 Click the **Restore** tab.

3 From the **Restore** tab, click **Recover replicated backups**.

The Recover wizard appears.




4 On the Destination page, select one of the following options:

- **Select a destination to use from an existing replication job**
- **Specify a new remote destination**

5 Click **Verify Authentication**.

6 On the Clients and Backups page, select a remote client and backups to recover back to this appliance. You can expand a client by clicking on the arrow to view its backups.

7 On the Ready to Complete page, review the items to recover.

8 Click **Finish** to start the recover request, or click **Back** to return to the previous screen if you need to make changes.

## Multi-Tenancy

VDP includes multi-tenancy support, which enables multiple customers or organizations to have separated accounts on a single VDP appliance. Each customer or organization can replicate to the VDP appliance as well as use the Replication Recovery feature to access their replicated data. Replicated data for a given customer or organization is only accessible using the customer or organization account credentials. Built-in credentials with full privileges such as root and repluser have access to all replicated data. These credentials should not be shared with individual customers or organizations. The replicated data for an account is isolated from the replicated data for all other accounts.

VDP, VDP Replication Target (available only in VDP 5.8), Avamar (version 7.1 or later recommended), and Avamar Virtual Edition (version 7.1 or later recommended) are the supported targets for multi-tenancy.

You can create multi-tenancy accounts by using the `create_av_domain.rb` script, which is provided on all VDP appliances. The `create_av_domain.rb` script defines the following parameters:

```
Usage: create_av_domain.rb

-c, --company=<Company-Name>       (Required)
-d, --department=<Department-Name> (Optional)
-u, --username=<User-Name>         (Required)
-p, --password=<User-Password>     (Required)
-h, --help
```

The company value should contain the name of the customer's company or organization. The optional department value allows for the setup of multiple accounts sharing the same company value but each having unique department values. The combined company and department values comprise an account. Each account has its own isolated bucket that can hold replicated data. The username and password parameters define the access credentials for an account. To create more than one set of access credentials for a single account, you can run the `create_av_domain.rb` script multiple times with the same values for company and department, but different values for username and password.

Follow these steps to execute the `create_av_domain.rb` script:

1 Log in (SSH) to the replication target server as the admin user.

2 Run the following command and provide the root password:

**su – root**

3 Change to the following directory:

**cd /usr/local/vdr/configure/bin**

4 Execute the following command and specify appropriate values for the account being created:

**./create_av_domain.rb --company=**<*Company-Name*> **--department=**<*Department-Name*> **--username=**<*User-Name*> **--password=**<*User-Password*>

for example:

**./create_av_domain.rb --company=Acme --department=Marketing --username=fred --password=topsecret**

After you create the account, enter the account information when you define a replication destination on the VDP appliance. The replication destination page requires a hostname or IP for the remote destination, the user credentials (username and password), and a path value. The path value is the account identifier, which is composed of the company and department values supplied when setting-up the account with the `create_av_domain.rb` script. A single forward-slash separates the company and department values. In the example above, the path value would be **Acme/Marketing**.

# Using File Level Restore 16

This chapter includes the following topics:


- "Introduction to the VDP Restore Client" on page 154
- "Logging in to the Restore Client" on page 155
- "Mounting Backups" on page 157
- "Filtering Backups" on page 157
- "Navigating Mounted Backups" on page 157
- "Performing File Level Restores" on page 157
- "Monitoring Restores" on page 159

# Introduction to the VDP Restore Client

VDP creates backups of entire virtual machines. These backups can be restored in their entirety using the VDP user interface through the vSphere web client. However, if you only want to restore specific files from these virtual machines, then use the VDP Restore Client (which is accessed through a web browser). This is called File Level Restore (FLR).

The Restore Client allows you to mount specific virtual machine backups as file systems, and then browse the file system to find the files you want to restore.

The Restore Client service is only available to virtual machines that have backups that are managed by VDP. This requires you to be logged in, either through the vCenter console or some other remote connection, to one of the virtual machines backed up by VDP.

**NOTE** To perform an FLR from the restore points that have been imported from the previously-used VDP disks, take at least one backup of the VM by using the new VDP.

**CAUTION** See "Software Requirements" on page 22 for web browsers supported by vSphere 6.0. Internet Explorer 10 is not supported, and unreliable with the restore client.

## LVM and EXT Support

Note the following when considering logical volumes managed by the Logical Volume Manager (LVM) and extended file systems:

- One physical volume (.vmdk) must be mapped exactly to one logical volume.
- Only ext2, and ext3 formatting (primary partition with master boot record (MBR) and standalone without MBR) are supported.
- LVM and ext4 are supported only with an external proxy.

## File Level Restore Limitations

FLR has the following limitations:

- VMware Tools must be installed on the target virtual machine. For best results, ensure that all virtual machines are running the latest available version of VMware Tools. Older versions are known to cause failures when browsing during the file-level restore operation.
- You cannot browse or restore the symbolic links.
- Browsing either a given directory contained within a backup or a restore destination is limited to a total of 5,000 files or folders.
- You cannot restore more than 5,000 folders or files in the same restore operation.
- When partitions are created, the lower ordered indices must be filled first. You cannot create a single partition and place it at the partition index 2, 3, or 4. The single partition must be at the partition index 1.
- FLR does not work if the VM is behind Network Address Translation (NAT).
- If the network appliance is using NAT, and sending all the VDP traffic through a proxy or a firewall, the FLR displays an error message that the client cannot be located in the VDP.
- FLR fails if you are using an older version of VMware tools. Ensure that you install the latest version of VMware tools on the target virtual machines.
- You cannot restore ACLs.
- You cannot log in to the Restore Client on a virtual machine if the virtual machine has been restored from a backup. The authentication fails because the Universal Unique Identifier (UUID) of the restored virtual machine is different.

  **Workaround**: Back up the new virtual machine and log in to the Restore Client by using the **Advanced Login** mode. "Advanced Login" on page 156 provides information.

- FLR does not work for Novell NSS volumes.

## Unsupported VMDK Configurations

FLR does not support the following virtual disk configurations:

- Unformatted disks
- Dynamic disks (Windows)
- GUID Partition Table (GPT) disks
- Extended partitions (Types: 05h, 0Fh, 85h, C5h, D5h)
- Virtual disk with more than one partition
- Two or more virtual disks mapped to a single partition
- Encrypted folders or files
- SCSI disks only are supported when restoring files or folders to the original virtual machine
- Zero-byte files
- Compressed partitions

**NOTE** In some cases (most notably extended partitions), you can restore the entire backup image to a temporary virtual machine, and then selectively copy the folders or files.

## Unsupported Windows Configurations

FLR does not support the following Windows 8 and Windows Server 2012 configurations:

- Deduplicated New Technology File System (NTFS)
- Resilient File System (ReFS)
- Extensible Firmware Interface (EFI) bootloader

# Logging in to the Restore Client

The Restore Client operates in either Basic or Advanced mode. You can only restore files from a Windows backup to a Windows machine, and you can only restore files from a Linux backup to a Linux machine.

**NOTE** If you are attempting to log in to Restore Client from a Windows 7 virtual machine, you must set the User Access Control (UAC) settings to the least restrictive settings for FLR to be able to function.

## Basic Login

With basic login, you connect to the Restore Client from a virtual machine that has been backed up by VDP. You log in to the Restore Client with the local administrative credentials of the virtual machine to which you are logged in, as shown in Figure 16-9.

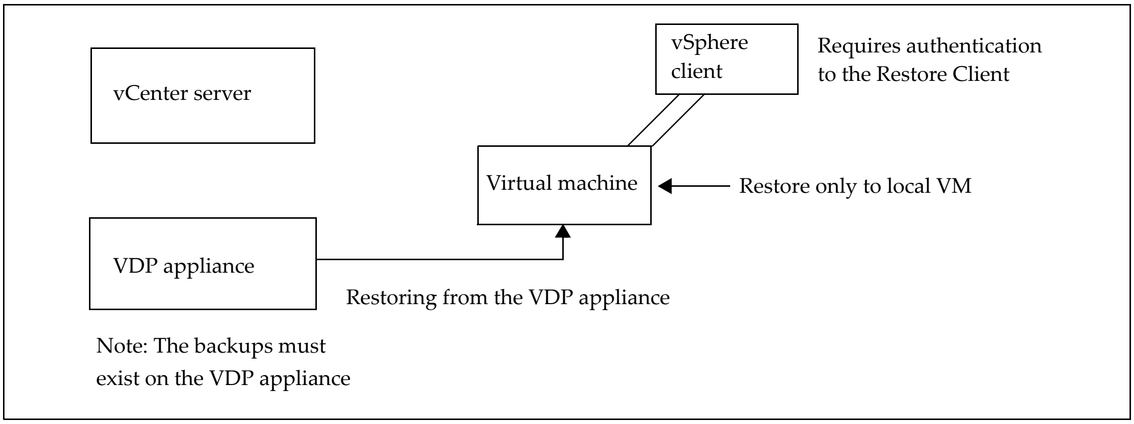



**Figure 16-9.** FLR Basic Login

For login instructions, refer to "Using the Restore Client in Basic Login Mode" on page 157.

With basic login, the Restore Client only displays backups for the local virtual machine. For example, if you are logged in to the Restore Client in basic mode from a Windows host named "WS44" then you are only able to mount and browse backups of "WS44."

## Advanced Login

With advanced login, you connect to the Restore Client from a virtual machine that has been backed up by VDP. You log in to the Restore Client with the local administrative credentials of the virtual machine that you are logged in to, and with the administrative credentials that were used to register the VDP appliance to the vCenter server as shown in Figure 16-10, or as the Active Directory user, inherit administrator user permissions from one of its groups.

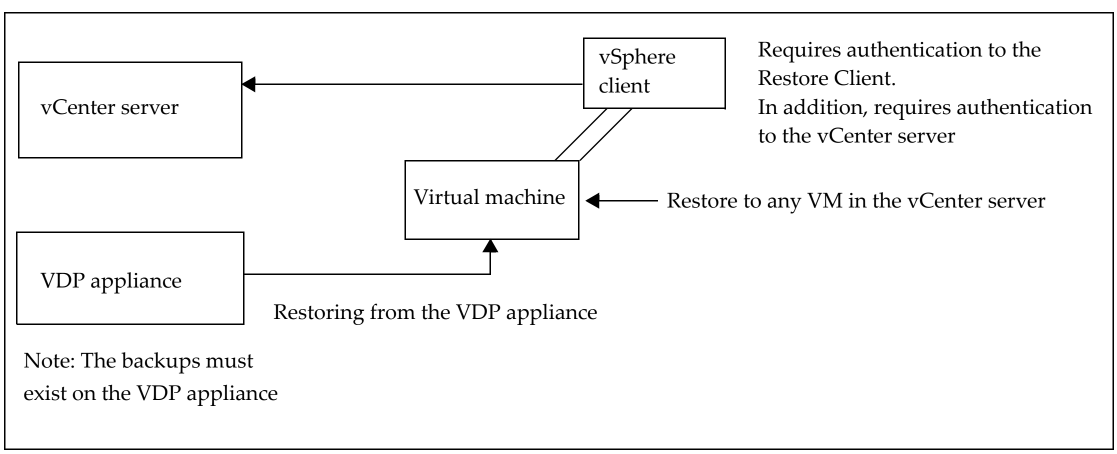



**Figure 16-10.** FLR Advanced Login

For login instructions, refer to "Using the Restore Client in Advanced Login Mode" on page 158.

After connecting to the Restore Client, you are able to mount, browse, and restore files from any virtual machine that has been backed up by VDP. All restore files are restored to the virtual machine to which you are currently logged in.

FLR advanced login requires you to use the same vCenter user credentials specified when the VDP appliance is installed. See "VDP Installation" on page 29 for additional information.

# Mounting Backups

After you successfully log in, the Manage mounted backups dialog appears, and by default, displays all the backups that are available to be mounted. The format of this dialog varies depending on how you log in.

- If you use basic login, the dialog lists all the backups from the client you logged into that can be mounted.
- If you use advanced login, the dialog lists of all clients that have backed up to VDP. Under each client, there is a list of all available backups to be mounted.

NOTE You can mount up to 254 vmdk file images by using the **Mount**, **Unmount**, or **Unmount** all buttons on the bottom right-hand corner of the dialog.

# Filtering Backups

In the Manage mounted backups dialog, you have the option of displaying all the backups or of filtering the list of backups. The list can be filtered in the following ways:

- **All restore points** display all backups.
- **Restore point date** display only backups within the specified date range.
- **VM name** displays only backups of hosts whose display name contains the text entered in the filter field. (This option is not available with basic login because only the backups belonging to the virtual machine you logged in with are displayed.)

# Navigating Mounted Backups

After backups have been mounted, you can navigate the contents of the backup by using the tree display on the left side of the Restore Client user interface. The appearance of the tree will vary depending on whether you used basic or advanced login.

# Performing File Level Restores

Using the main screen of the Restore Client, you can restore specific files by navigating the file system tree in the left-hand column, and then clicking directories in the tree or clicking files or directories in the right-hand column.

## Using the Restore Client in Basic Login Mode

Use the Restore Client on a Windows or Linux virtual machine in basic login mode to access individual files from restore points for that machine, rather than restoring the entire virtual machine.

### Prerequisites

- Verify that vSphere Data Protection (VDP) is installed and configured on your vCenter server.
- For basic login, you can only log in to the Restore Client from a virtual machine that has been backed up by VDP.
- VMware Tools must be installed on the virtual machine in order to perform file-level restores from backups. Refer to the VMware website for list of operating systems that support VMware Tools.

### Procedure

1 Use Remote Desktop or a vSphere web client to access the local host that has been backed up through VDP.

2 Access the VDP Restore Client:

**https://**<*IP_address_of _VDP_appliance*>**:8543/flr**

3 On the **Credentials** page under **Local Credentials**, specify the **Username** and **Password** for the local host and click **Login**.

The **Manage Mounted Backups** dialog box appears and lists all of the restore points for the client you are accessing.

4 Select the mount point to restore and click **Mount**.

After the mounting completes, the drive icon will appear as a green networked drive  .

5 Click **Close**.

6 In the **Mounted Backups** window, navigate to and select the folders and files to restore.

7 Click **Restore selected files.**

8 In the **Select Destination** dialog box, navigate to and select the drive and the destination folder to restore.

9 Click **Restore**.

An Initiate Restore confirmation dialog box appears.

10 Click **Yes**.

A successfully initiated dialog box appears.

11 Click **OK**.

12 Select the **Monitor Restores** tab to view the status of the restore and ensure that the restore succeeds.

## Using the Restore Client in Advanced Login Mode

Use the restore client on a Windows or Linux virtual machine in advanced login mode to access virtual machines on a vCenter server that contain restore points to perform file-level recovery.

### Prerequisites

- Verify that VDP is installed and configured on your vCenter server.
- FLR advanced login requires you to use the same vCenter user credentials specified when the VDP appliance is installed. See "VDP Installation" on page 29 for additional information.
- VMware Tools must be installed on the virtual machine in order to perform file-level restores from backups. Refer to the VMware website for list of operating systems that support VMware Tools.

### Procedure

1 Log in remotely using Remote Desktop, or use a vSphere web client to access a virtual machine.

2 Access the VDP Restore Client:

**https://**<*IP_address_of _VDP_appliance*>**:8543/flr**

3 On the **Credentials** page, under **Local Credentials**, specify the **Username** and **Password** for the local host. In the **vCenter Credentials** field, specify the vCenter administrator **Username** and **Password** and click **Login**.

The **Manage Mounted Backups** dialog box appears. It lists all the restore points for all the clients of the same type that are backed up to VDP.

4 Select the mount point to restore and click **Mount**.

After the mounting completes, the drive icon will display as a green networked drive .

5 Click **Close**.

6 In the **Mounted Backups** window, navigate to and select the virtual machine, folders, and files for recovery.

7 Click **Restore selected files.**

8 In the **Select Destination** dialog box, navigate to and select the drive and destination folder for recovery.

9 Click **Restore**.

An Initiate Restore confirmation dialog box appears.

10 Click **Yes**.

A successfully initiated dialog box appears.

11 Click **OK**.

12 Select the **Monitor Restores** tab to view the status of the restore and ensure that the restore succeeds.

## Monitoring Restores

To monitor current and past activity of the Restore Client, click the **Monitor Restores** button. The monitor restore screen displays information about current and recently-completed restore operations.

The columns in this table are sortable by clicking on the column heading. Clicking multiple times on a table heading will reverse the sort order, and an up or down arrow reflects whether the sort order is ascending or descending.

By default, **Monitor Restores** shows all the jobs that in are either in progress or that have completed during the current session. If you want to see jobs that completed or failed in a previous session, select **Show Completed Activities**. All the previous completed and failed jobs are displayed along with running and pending jobs.

# 17 VDP Application Support

This chapter includes the following topics:


- "VDP Application Support" on page 162
- "Backing Up and Restoring Microsoft SQL Servers" on page 163
- "Backing Up and Restoring Microsoft Exchange Servers" on page 172
- "Backing Up and Restoring Microsoft SharePoint Servers" on page 185

# VDP Application Support

VDP supports granular guest-level backup and recovery support for Microsoft Exchange Servers, SQL Servers, and SharePoint Servers. To support guest-level backups, a VDP client is installed on the Exchange Servers, SharePoint Servers, and SQL Servers.

## Installing Application Agents

To install application agents, refer to the following application-specific instructions:


- "Installing VDP for SQL Server Client" on page 165
- "Installing VDP for Exchange Server Client" on page 174
- "Installing VDP for SharePoint Server Client" on page 186


## Checking the User Account Control Setting on Microsoft Windows

The User Account Control (UAC) feature limits application software to standard user privileges. You must provide administrator privileges for certain tasks, such as installing software. UAC is enabled by default.

If you start an VDP client or plug-in installer without administrator privileges on a computer with UAC enabled, the software does not install correctly. You can disable or bypass UAC. The installation procedures in this chapter provide one method to bypass UAC. Other methods and additional information are available in the Microsoft documentation.

## Installing VDP Clients when UAC is Enabled

When you attempt to install VDP clients with UAC enabled, the installation displays the following error:

The VMware VDP for <*Microsoft Application*> Server cannot be installed. Please confirm that you are logged in as an administrator and that all installation prerequisites have been met.

To resolve this issue, you must run the installer with administrative privileges by performing the following steps:

1. In Windows, right-click the Command Prompt icon and select **Run as administrator**.
2. In the Command Prompt window, change the working directory to the location of the installation package by typing the following path:

   **cd** *install_path*

   where *install_path* is the full path of the temporary folder that contains the installation package.
3. Type the appropriate command to start the installer:

   **msiexec /i VMwareVDPExchange-windows-x86_64-**<*version*>**.msi**
   **msiexec /i VMwareVDPMoss-windows-x86_64-**<*version*>**.msi**
   **msiexec /i VMwareVDPSQL-windows-x86_64-**<*version*>**.msi**
   **msiexec /i VMwareVDPSQL-windows-x86_32-**<*version*>**.msi**

   where *version* is the VDP client version.

# Backing Up and Restoring Microsoft SQL Servers

VDP supports enhanced backup and restore options for Microsoft SQL Servers.

This section covers the following topics:


- "Microsoft SQL Server Options" on page 163
- "Hardware Requirements" on page 163
- "Authentication Requirements" on page 164
- "Microsoft SQL Server Support" on page 164
- "Installing VDP for SQL Server Client" on page 165
- "Configuring the Cluster Client in a Failover Cluster" on page 166
- "Configuring the Cluster Client for an AlwaysOn Availability Group" on page 167
- "Creating Backup Jobs for Microsoft SQL Servers" on page 168
- "Restoring Backups of Microsoft SQL Servers" on page 171
- "Monitoring Client Activity" on page 172
- "Uninstalling the VDP Plug-in for SQL Server" on page 172


## Microsoft SQL Server Options

The following options are supported for Microsoft SQL Servers:

- Backup selected SQL Servers
- Select entire database instances for backup
- Select individual databases for backup
- Support for full, differential, or incremental backups
- Support to use incremental backups after full backups
- Support for multi-streaming backups (up to six streams)
- Support for simple-mode database backups (skips incremental)
- Restore to original or alternate location
- Restore a database in the original instance using specified path
- Restore a database to a different instance using specified path

## Hardware Requirements

Table 17-24 lists the hardware requirements for the Microsoft SQL Server.

**Table 17-24.** Hardware Requirements for Microsoft SQL Server

| Requirement | Minimum |
|---|---|
| Memory (RAM) | 512 MB (2 GB recommended) |
| Hard drive space | 1 GB permanent hard drive space for software installation.<br>The Microsoft SQL Server software also requires an additional 12 MB of permanent hard drive space for each 64 MB of physical RAM. The space is necessary for local cache files. |

## Authentication Requirements

Avamar connects to SQL Server for backups and restores by using one of the following methods:

- Windows authentication, that is, the Windows system service account (NT AUTHORITY\SYSTEM), which has privileges on the SQL Server. The account must have the sysadmin server-level role.
- SQL Server authentication, that is, the SQL Server login account. The account must have the sysadmin server-level role. When you configure the SQL Server, select the mixed authentication mode for the database engine.

## Microsoft SQL Server Support

VDP supports the following versions of the SQL Server:

- SQL Server failover clusters for the following SQL versions:
  - SQL Server 2014
  - SQL Server 2012
  - SQL Server 2008, 2008 R2
  - SQL Server 2005
- SQL AlwaysOn clusters for the following SQL versions:
  - SQL Server 2014
  - SQL Server 2012
- SQL Server 2014
  - SQL Server 2014 (x86/x64) on Windows Server 2012
  - SQL Server 2014 (x86/x64) on Windows Server 2008 SP2 or later
  - SQL Server 2014 (x86/x64) on Windows Server 2008 R2 SP1 or later
- SQL Server 2012 (x86/x64) on Windows Server 2012
- SQL Server 2012 (x86) on Windows Server 2008 SP2 or later
- SQL Server 2012 (x64) on Windows Server 2008 R2 SP1 or later
- SQL Server 2008 R2 and later on:
  - Windows Server 2008 or later (x86/x64)
  - Windows Server 2008 R2 (x64)
  - Windows Server 2012
- SQL Server 2008 SP1 or later on:
  - Windows Server 2008 SP1 or later (x86/x64)
  - Windows Server 2008 R2 (x64)
  - Windows Server 2012
- SQL Server 2005 SP3 on:
  - Windows Server 2008 SP1 or later (x86/x64)
  - Windows Server 2008 R2 (x64)
- SQL Server 2005 SP2 on Windows Server 2008 (X86/x64)

## Installing VDP for SQL Server Client

To support guest-level backups, the VDP for SQL Server Client must be installed on each SQL Server for backup and restore support.

To install the VDP for SQL Server Client in a cluster, install the VDP for SQL Server Client on each node, register each node, and then configure the VDP cluster client. To install the VDP for SQL Server Client in a cluster, perform the following steps:

1. Install the VDP for SQL Server Client in the same folder on each node in the cluster.

   The installation process installs the software, and then registers and activates each node in the cluster with the VDP appliance.

2. Use the VMware VDP Windows Cluster Configuration Wizard to configure the VDP appliance.

### Prerequisites

- Before using VDP, you must install and configure the VDP appliance described in "VDP Installation and Configuration" on page 21 and you must have administrative rights to the SQL Server.
- The following software must be installed on the SQL Server:
  - .NET 4.0
  - SQL Server Installation Component
  - Client Tools SDK

### Procedure

1. On each SQL Server client, access the vSphere web client:

   **https://**<*IP_address_vCenter_Server*>**:9443/vsphere-client/**

2. In the Credentials page, enter an administrative vCenter username and password and click **Login**.
3. In the vSphere web client, select **VDP**.
4. In the Welcome to VDP page, select the VDP appliance and click **Connect**.
5. Click the **Configuration** tab.
6. In **Client Downloads**, click **Microsoft SQL Server 32-bit** or **Microsoft SQL Server 64-bit** (based on the version of the SQL Server client).
7. Depending on your browser, you can save `.msi` file or you can run it. When you run the `.msi` file, the VMware VDP for SQL Server Setup wizard starts. Click **Next**.
8. On the End-User License Agreement page, read the license and if acceptable, click **I accept the terms in the License Agreement**, and click **Next**.
9. On the appliance Registration Information page, type the name of the VDP appliance that will back up the SQL Server, and click **Next**.
10. On the Ready to install VMware VDP for SQL Server page, click **Install**.
11. On the Completed the VMware VDP for SQL Server Setup Wizard page, click **Finish**.

Repeat this procedure for additional SQL Servers.

## Configuring the Cluster Client in a Failover Cluster

The VDP cluster client in a failover cluster enables you to back up and restore SQL Server data on shared storage in the cluster, regardless of which node is managing the data at the time of the backup or restore. The VMware VDP Windows Cluster Configuration Wizard provides the steps to configure the cluster client for the SQL Server plug-in in a failover cluster.

### Procedure

1 Log in to the active node in the cluster as a domain administrator. The account must also be a member of the local Administrators group on each cluster node.

2 Start the VMware VDP Windows Cluster Configuration Wizard:

- On Windows Server 2012, open the **Start** screen and select **VMware VDP Windows Cluster Configuration Wizard.**
- On Windows Server 2008, open the **Start** menu and select **Program Files > VMware VDP > VMware VDP Windows Cluster Configuration Wizard**.

3 On the **Welcome** page, click **Next**.

4 On the **Plug-Ins** page, select **SQL** and click **Next**.

The Cluster Nodes page appears with a list of nodes and their status.

5 On the **Cluster Nodes** page, ensure that the environment meets the following requirements and click **Next**:

- The status for each SQL Server node is Up.
- The installation status of the Windows client software for each node is Installed.
- The installation status of the SQL Server plug-in on each node is Installed.

6 On the **Operations** page, select **Configure a new cluster client for all nodes** and click **Next**.

7 On the **Prerequisites** page, ensure that the environment meets all prerequisites.

A check mark next to a prerequisite indicates that the environment meets the prerequisite.

If the environment does not meet a prerequisite, exit the wizard, resolve the issue, and then restart the wizard.

8 Select the IP version that the environment uses and click **Next**.

9 On the **SQL Settings** page:

a Select the cluster group, service, or role for the cluster client from the list:

- On Windows Server 2012, select the cluster role for the cluster client from the **Cluster role for cluster client** list.
- On Windows Server 2008, select the cluster service for the cluster client from the **Cluster service for cluster client** list.

b Select the shared volume for the cluster client from the Shared volume for the cluster client list.

c Click **Next**.

10 On the **Server Settings** page, specify the settings for the VDP appliance:

a Type either the DNS name of the VDP appliance in the **Name** box or the IP address in the **IPv4/IPv6 address** box.

b Type the data port for VDP client/server communication in the **Port number** box.

**NOTE** Port 28001 is the default port that the VDP client uses to communicate with the VDP appliance.

c Type the name of the folder or volume in the **Cluster client's var directory** box, or click **Browse** to select a folder or volume.

This folder or volume stores the cluster client configuration and log files. All nodes in the cluster must have write access to this folder or volume.

**NOTE** Select a volume that the cluster owns instead of a remote pathname on the network.

d Click **Next**.

11 On the **Summary** page, review the configuration settings and click **Configure**.

The **Progress** page provides the status of the configuration.

12 On the **Results** page, review the information and click **Close**.

## Configuring the Cluster Client for an AlwaysOn Availability Group

The VDP cluster client for an AlwaysOn availability group enables you to back up and restore SQL Server databases in an availability group. The VMware VDP Windows Cluster Configuration Wizard provides the steps to configure the VDP cluster client for the SQL Server plug-in in an AlwaysOn availability group.

### Procedure

1 Log in to a cluster node as a domain administrator. The account must also be a member of the local Administrators group on each cluster node.

2 Start the VMware VDP Windows Cluster Configuration Wizard:

- On Windows Server 2012, open the **Start** screen and select **VMware VDP Windows Cluster Configuration Wizard**.
- On Windows Server 2008, open the **Start** menu and select **Program Files > VMware VDP > VMware VDP Windows Cluster Configuration Wizard**.

3 On the **Welcome** page, click **Next**.

4 On the **Plug-Ins** page, select **SQL AlwaysOn** and click **Next**.

5 On the **Cluster Nodes** page:

a Ensure that the environment meets the following requirements:

- The status for each SQL Server node is Up.
- The installation status of the Windows client software for each node is Installed.
- The installation status of the SQL Server plug-in on each node is Installed.

b Click **Next**.

6 On the **Operations** page, select **Configure a new cluster client for all nodes** and click **Next**.

7 On the **Prerequisites** page, ensure that the environment meets all prerequisites.

A check mark next to a prerequisite indicates that the environment meets the prerequisite.

If the environment does not meet a prerequisite, then exit the wizard, resolve the issue, and restart the wizard.

8 Select the IP version that the environment uses and click **Next**.

9 On the **SQL AlwaysOn Settings** page:

a Select the cluster group, service, or role for the cluster client from the list:

- On Windows Server 2012, select the cluster role for the cluster client from the Cluster role for cluster client list.
- On Windows Server 2008, select the cluster service for the cluster client from the Cluster service for cluster client list.

The name of the cluster client appears in the **Cluster client name** field.

**NOTE** You must configure an availability group listener for each availability group. Do not configure a cluster client for an availability group that does not have a listener.

   b Click **Next**.

10 On the **Server Settings** page, specify the settings for the VDP appliance:

   a Type either the DNS name of the VDP appliance in the **Name** box or the IP address in the **IPv4/IPv6 address** box.

   b Type the data port for VDP client/server communication in the **Port number** box.

**NOTE** Port 28001 is the default port that the VDP appliance uses to communicate.

   c Type the name of the folder or volume in the **Cluster client's var directory** box, or click **Browse** to select a folder or volume.

   This folder or volume stores the cluster client configuration and log files. All nodes in the cluster must have write access to this folder or volume.

**NOTE** Select a volume that the cluster owns instead of a remote pathname on the network.

   d Click **Next**.

11 On the **Summary** page, review the configuration settings and click **Configure**.

   The **Progress** page provides the status of the configuration.

12 On the **Results** page, review the information and click **Close**.

## Creating Backup Jobs for Microsoft SQL Servers

The VMware VDP for SQL Server Client must be installed on each SQL Server that will be available for backup. See "Installing VDP for SQL Server Client" on page 165 for additional information on client installation.

1 On the vSphere web client, select the **Backup** tab.

2 In the **Backup** tab, click **Backup Job Actions** and select **New** to start the **Create a new backup job** wizard.

   On the Job Type page of the wizard, select **Applications**. This option lets you back up the full server or selected databases.

### Backing up Applications

If you select **Applications** on the Job Type page, you can choose to back up application servers or individual databases.

1 On the **Job Type** page of the **Create a new backup job** wizard, select **Applications** and click **Next**.

2 On the **Data Type** page, select one of the following options and click **Next**:

   - **Full Server**: This option lets you back up entire applications servers.
   - **Selected Databases**: This option enables you to back up individual application server databases.

3 On the **Backup Sources** page, click the arrow beside one of the following backup sources to expand the list:

   - **Microsoft SQL Server(s)**: Select for SQL Server backup.
   - **Microsoft SQL Failover Cluster(s)**: Select for SQL Failover Cluster backup.
   - **Microsoft SQL AlwaysOn Cluster(s)**: Select for SQL AlwaysOn availability group backup.

4 Perform one of the following steps:

   - If you chose to back up a full server, select the checkbox next to the SQL server that you want to back up, and then click **Next**.

     **NOTE** Best practice is to select only one SQL Server per backup job.

- If you chose to back up selected databases, click the arrow beside an SQL Server, continue drilling down until you can select the database or storage group that you want to back up, and click **Next**.

5 On the **Backup Options** page, select a backup type of **Full**, **Differential**, or **Incremental**. The options that you can configure depend on which one you select.

- **Full**: The **Full** option backs up the entire database, including all objects, system tables, and data. The options for full backup are described as follows:

  - **Force incremental backup after full backup**: Selecting or clearing this checkbox specifies whether to force an incremental backup that contains the transactions that occur between full backups. This creates a point-in-time recovery to a point between full backups.

    Do not use this option on databases that use the simple recovery model because those databases do not support transaction log backups. This includes system databases such as the master and msdb databases.

    For simple model recovery databases, use the **For simple model recovery databases** option.

  - **Enable multi-stream backup**: You can either back up multiple databases in parallel with one stream per database, or back up a single database using multiple parallel streams. If you choose to back up a single database with multiple parallel streams, you can specify the minimum size of each stream during the backup.

    After you determine the minimum stream size, you can calculate the number of streams used to back up the database using the following equation:

    Database size/minimum stream size = Number of streams

    For example, if a database is 1,280 MB and you set the minimum stream size to the default setting of 256 MB, then the number of streams that are used to perform a full backup of the database is five, as shown in the following equation:

    1,280 MB/256 = 5

    For transaction log and differential backups, the size of the data to back up, and not the total database size, is used to calculate the number of streams. If the database size is less than the minimum stream size, VDP uses a single stream to back up the database.

    If you calculate the number of streams for a database based on the minimum stream size, and the number exceeds the maximum number of streams that you configured for the backup, the backup of the database uses only the maximum number of streams.

  - **For simple recovery model databases**: This option specifies how VDP handles incremental (transaction log) backups of databases that use the simple recovery model, which does not support transaction log backups:

    **Skip incremental with error** (default setting): If you select databases with different recovery models for the backup, the backup does not include the databases with the simple recovery model. The backup completes with exceptions, and an error message is written to the log. If you select only databases with the simple recovery model for the backup, the backup fails.

    **Skip incremental with warning:** If you select databases with different recovery models for the backup, the backup does not include databases with the simple recovery model. The backup completes successfully, and a warning is written to the log for each database that uses the simple recovery model. If you select only databases with the simple recovery model for the backup, the backup fails.

    **Promote incremental to full**: A full backup occurs automatically instead of a transaction log backup for databases that use the simple recovery model.

- **Truncate database log**: This option specifies how database transaction log truncation behavior is controlled. Truncate options include the following:

  **Only for incremental backup** (default setting): The database transaction log is truncated if the backup type is set to incremental (transaction log). No log truncation occurs if the backup type is full or differential.

  **For all backup types**: The database transaction log is truncated regardless of the backup type. This setting breaks the chain of log backups and should not be used unless the backup type is set to full.

  **Never**: The database transaction log is not truncated under any circumstances.

- **Authentication method:** The authentication method specifies whether to use NT authentication or SQL Server authentication to connect to SQL Server. If you select SQL Server authentication, specify the SQL Server login name and password.

- **Availability group replica for backup:** There are four options:

  **Primary:** If selected, the backup is executed on the primary replica of selected AlwaysOn availability group.

  **Prefer secondary:** If selected, the backup is executed on the secondary replica of selected AlwaysOn availability group. If no secondary replica is available, the backup will be performed on the primary replica.

  **Secondary only:** If selected, the backup is executed on the secondary replica of selected AlwaysOn availability group. If no secondary replica is available, the backup will be interrupted and an appropriate error message will be written to the log file.

  **SQL Server defined:** If selected, the backup is executed on the primary or secondary replica based on the SQL server configuration. If Automated_Backup_Preference is set to none the backup will be executed on primary replica.

- **Differential** or **Incremental**: The **Differential** option backs up any data that has changed since the last full backup. The **Incremental** option backs up only the transaction logs. The only configuration option that differs from a full backup is that you can force a full backup rather than an incremental backup.

  - **Force full backup**: Selecting or clearing the checkbox determines whether to perform a full backup when VDP detects a log gap or when there is no previous full backup, from which a transaction log (incremental) or differential backup can be applied. Effectively, this option automates taking a full backup when necessary.

    If you select **Differential** or **Incremental**, you should leave this option selected (the default setting). Otherwise, you might not be able to restore data in the event that no existing full backup is present on VDP.

6. Click **Next**.
7. On the **Schedule** page, select the backup schedule and start time for the backup job and click **Next**.

   "Specifying the Backup Schedule" on page 115 provides additional information about how to configure the schedule.
8. On the **Retention Policy** page, select an option for how long to keep the backup and click **Next**.

   "Setting the Retention Policy" on page 115 provides additional information about how to configure the retention policy.
9. On the **Name** page, type a name for the backup job and click **Next**.
10. On the **Ready to Complete** page, review the summary information about the backup job and click **Finish**.
11. Click **OK** when you see the confirmation that the backup job was created successfully.

## Restoring Backups of Microsoft SQL Servers

Once backups have been run on Microsoft SQL Servers, you can restore these backups to their original location or to an alternate location.

### Procedure

1. In the vSphere web client, select the **Restore** tab.
2. Select the backup that you want to restore. While you can select multiple SQL Servers, you can only select one restore point for each SQL Server.
3. Click **Restore**.
4. On the **Select Backup** page, select the backup job that you want to restore and click **Next**.
5. On the **Select Restore Options** page:

   a. Perform one of the following steps:

      - Leave the **Restore to Original Location** option selected (the default setting) to restore the backup to its original location.
      - Clear the **Restore to Original Location** option to restore the backup to an alternate location, and then do the following:

        i. Click **Choose** to select the destination client.

        ii. In the **SQL instance** box, type the name of the SQL instance. If you use “local,” it must be in parentheses.

        iii. In the **Location path** box, type the existing full Windows path where the database files will be restored.

        If the Location path does not exist, it will not be created, and the restore will fail.

        iv. In the **Log file path** box, type the existing full Windows path where the log files will be restored.

   b. If you want to specify advanced options, click the arrow beside **Advanced options** to expand the list. The options are described as follows:

      - **Use SQL REPLACE:** This option specifies that the SQL Server should create any necessary database and related files even if another database or file already exists with the same name.

        This option overrides a SQL Server safety check that is intended to prevent you from accidentally overwriting a different database or file. This safety check is described in the Microsoft Transact-SQL Reference Manual under the RESTORE command section.

      - **Restore only on primary replica**: Sets the `--recover-primary-only` flag, which disables the automatic recovery of secondary replicas and performs the recovery only on the primary replica. By default, the check box is enabled, so the recovery will be performed only on the primary replica. This option is enabled only for databases that reside in an AlwaysOn availability group.

**NOTE** After you restore a database on only the primary replica, the corresponding database on the secondary replica is in a restoring state.

      - **Tail-log backup:** To perform a tail-log backup during the restore process, the database must be online and using either the full or bulk-logged recovery model. You cannot perform a tail-log backup on the system databases because those databases use the simple recovery model (for example, the master and msdb databases).

        Do not select tail-log backup if you are performing a redirected restore to a different SQL Server instance.

- **Restore system databases**: You rarely need to restore only system databases. However, the restore might be required if one or more system databases are damaged.

  You will more likely need to restore system databases at the same time that you restore user databases. When you select both the system and user databases for restore, the system databases are restored first.

  When you restore system databases, the VDP Microsoft SQL Server Client automatically restores the databases in the correct order (master, msdb, then model) managing SQL Server services.

- **Authentication method**: The authentication method specifies whether to use NT authentication or SQL Server authentication to connect to SQL Server. If you select SQL server authentication, specify the SQL Server login name and password.

- **Recovery operation**: Select the appropriate value—**RECOVERY** or **NORECOVERY** if you are performing the restore in a SQL AlwaysOn cluster environment.

c Click **Next**.

6 On the **Ready to complete** page, review the restore requests and click **Finish**.

7 In the message box that states that the restore has been successfully initiated, click **OK**.

8 Monitor the restore's progress in the **Recent Tasks** panel.

### Monitoring Client Activity

You can monitor tasks and events for all the client activity by collecting and analyzing the client logs. The client logs are Microsoft application (MSApp)-related logs. The aggregated client log includes any replication, backup, restore, or automatic backup verification (ABV) job that passed with exceptions or failed. For more information, refer to "Collecting VDP Logs or Diagnostic Information" on page 40.

### Uninstalling the VDP Plug-in for SQL Server

To uninstall the VDP Plug-in for SQL Server:

- On Windows Server 2012 or Windows Server 2008, use **Programs and Features**.
- On Windows Server 2003, use **Add/Remove Programs**.

## Backing Up and Restoring Microsoft Exchange Servers

This section covers the following topics:


- "Microsoft Exchange Server Support" on page 173
- "Installing VDP for Exchange Server Client" on page 174
- "Using the VMware Exchange Backup User Configuration Tool" on page 176
- "Manually Configuring the VDP Backup Service" on page 178
- "Creating Backup Jobs for Microsoft Exchange Servers" on page 178
- "Restoring Backups of Microsoft Exchange Servers" on page 180
- "Granular Level Recovery on Microsoft Exchange Servers" on page 182
- "Uninstalling the Exchange Server Plug-in" on page 181

## Microsoft Exchange Server Options

VDP supports the enhanced backup and restore options for Microsoft Exchange Servers:

- Backup selected Exchange Servers
- Backup selected individual Exchange databases or storage groups
- Ability to perform incremental backups
- Support for multi-streaming backups (up to ten streams)
- Support for circular logging (promote, circular, and skip)
- Ability to restore Exchange to original location or alternate location
- Option for no replay logs during restore
- Recovery Storage Groups (RSG)/Recovery database (RDB)
- Granular-level restores

## Microsoft Exchange Server Support

Table 17-25 lists the Microsoft Exchange Server versions and operating systems that the VDP plug-in for Microsoft Exchange supports.

**Table 17-25.** Supported Microsoft Exchange Server Versions and Operating Systems

| Exchange Server version | Operating systems |
|---|---|
| ■ Exchange Server 2013<br>■ Exchange Server 2013 database availability group (DAG) | ■ Windows Server 2012 (x64)<br>■ Windows Server 2012 R2 (x64)<br>■ Windows Server 2008 R2 (x64) |
| ■ Exchange Server 2010 SP3<br>■ Exchange Server 2010 Database Availability Group (DAG) | ■ Windows Server 2012 (x64)<br>■ Windows Server 2008 R2 (x64)<br>■ Windows Server 2008 SP2 (x64) |
| ■ Exchange Server 2007 SP3 | ■ Windows Server 2008 R2 (x64)<br>■ Windows Server 2008 SP2 (x64) |

## Microsoft .NET Framework 4 Requirement

The Exchange Server VSS plug-in requires installation of Microsoft .NET Framework 4 on each server in the Exchange Server forest. Search the Microsoft Download Center for "Microsoft .NET Framework 4" to find downloads and additional information.

## Hardware Requirements

Table 17-26 lists the hardware requirements for the VDP Plug-in for Microsoft Exchange Server.

**Table 17-26.** Hardware Requirements for Microsoft Exchange Server

| Requirement | Minimum |
|---|---|
| Memory (RAM) | 64 MB |
| Hard drive space | Software installation requires at least 100 MB of permanent hard drive space, with 1 GB recommended.<br>Local cache files require an additional 12 MB of permanent hard drive space for each 64 MB of physical RAM. |

## Unsupported Microsoft Exchange Servers

Microsoft Exchange Server 2007 clusters (SCC, CCR, SCR) are not supported with the VDP plug-in for Microsoft Exchange Server.

## Installing VDP for Exchange Server Client

To support guest-level backups, the VMware vSphere Data Protection (VDP) for Exchange Server Client must be installed on each Exchange Server for backup and restore support.

### Prerequisites

Before using VDP, you must install and configure the VDP appliance described in "VDP Installation and Configuration" on page 21 and you must have administrative rights to the Exchange Server.

### Procedure

1. On each Exchange Server client, access the vSphere web client:

   **https://**<*IP_address_vCenter_Server*>**:9443/vsphere-client/**

2. In the Credentials page, enter an administrative vCenter username and password, and then click **Login**.
3. On the vSphere web client, select **VDP**.
4. In the Welcome to VDP page, select the VDP appliance and click **Connect**.
5. Click the **Configuration** tab.
6. In **Client Downloads**, click **Microsoft Exchange Server 64-bit**.
7. Depending on your browser, you can save the `.msi` file or run it. When you run the `.msi` file, the VMware VDP for Exchange Server Setup wizard starts. Click **Next**.
8. On the **End-User License Agreement** page, read the license and if acceptable, click **I accept the terms in the License Agreement** and click **Next**.
9. On the **Appliance Registration Information** page, type in the IP address or fully qualified domain name of the VDP appliance that will back up the Exchange Server and click **Next**.
10. (Optional) Select the option to install the **Exchange Server GLR** plug-in if you plan to use the server for granular level recovery.

NOTE You must reboot the Microsoft Exchange Server if you select the Exchange Server GLR option.

11. On the **Ready to install VMware VDP for Exchange Server** page, click **Install**.
12. On the **Completed the VMware VDP for Exchange Server Setup Wizard** page, click **Finish**.

If you selected the **VDP Exchange Backup User Configuration Tool** checkbox, proceed to "Using the VMware Exchange Backup User Configuration Tool" on page 176.

If you did not select the **VDP Exchange Backup Configuration Tool** checkbox, proceed to "Manually Configuring the VDP Backup Service" on page 178.

Repeat this procedure for additional Exchange Servers.

## Installing in a DAG or Cluster Environment

### Procedure

1. Exchange for VDP consists of a single, standalone installer. Install the VDP for Exchange Client Plug-in on each Microsoft Exchange server with the Mailbox server role.

   To use a server for granular level recovery, select the options to install both the Exchange GLR plug-in and the Exchange VSS plug-in. In a DAG environment, you should configure at least one server for GLR.

NOTE You must restart the Exchange server after you install the Exchange GLR plug-in.

2. Register each Exchange server as a client with the VDP appliance.
3. Create and configure the VmwareVDPBackupUser account.

4 Use the VMware VDP Windows Cluster Configuration Wizard to configure the Exchange DAG client or the VDP cluster client.

## Configuring an Exchange DAG Client

Configuring an Exchange DAG client allows you to perform federated backups of databases in an Exchange Server 2013 or 2010 DAG.

### Prerequisites

Verify that the DAG cluster environment meets the following prerequisites before configuring the Exchange DAG client.

- The VDP Windows Client has been installed.
- The VDP Backup Plug-in for Exchange DAG has been installed.
- The DAG Group exists if the DAG client is already configured for any nodes.
- An unused static IP address is available to be assigned to the new VDP Exchange DAG Client.
- The machine accounts for all the cluster nodes must have full access to the SMB share.
- The Exchange DAG federated backup must have a network share created (to be used as the `var` directory).

### Procedure

1. Log in to an Exchange server in the DAG with the VmwareVDPBackupUser account.
2. Start the VMware VDP Windows Cluster Configuration Wizard:
   - On Windows Server 2012, open the **Start** screen and select **VMware VDP Windows Cluster Configuration Wizard**.
   - On Windows Server 2008, open the **Start** menu and select **Program Files > VMware VDP > VMware VDP Windows Cluster Configuration Wizard**.
3. On the **Welcome** page, click **Next**.
4. On the **Plug-Ins** page, select **Exchange DAG** and click **Next**.
5. On the **DAG Nodes** page:
   a. Ensure that the environment meets the following requirements:
      - The status for each Exchange server is Up.
      - The installation status of the Windows client software for each server is Installed.
      - The installation status of the Exchange VSS plug-in on each server is Installed.

   b. Click **Next**.
6. On the **Operations** page, select **Configure a new DAG client for all nodes** and click **Next**.
7. On the **Prerequisites** page:
   a. Ensure that the environment meets all prerequisites.

      A check mark next to a prerequisite indicates that the environment meets the prerequisite.

      If the environment does not meet a prerequisite, then exit the wizard, resolve the issue, and restart the wizard.

   b. Select the Internet Protocol version—IPv4 or IPv6 that the environment uses.

   c. Click **Next**.

8 On the **DAG Client Settings** page, specify the client settings for the cluster group for the DAG client:

   a Select the network in the network list.

   b Type the IP address for the DAG client cluster group in the **Exchange DAG client IPv4/IPv6 address** box. The IP address must be a unique, unused IP address. Do not use the IP address for the DAG.

   c Type the network mask for the DAG client cluster group in the **Exchange DAG client IP subnet mask** box.

   d Click **Next**.

9 On the User Settings page:

   a Select one of the following login accounts:

      - Local System account
      - This account, and specify the Account Name and Password for the VmwareVDPBackupUser account.

   b Click **Next**.

10 On the **Server Settings** page, specify the settings for the VDP appliance:

   a Type either the DNS name of the VDP appliance in the **Name** box or the IP address in the **IPv4/IPv6 address** box.

   b Type the name of the VDP domain for the Exchange DAG client in the **VDP client domain for the DAG client** box.

   c Type the data port for the VDP client/server communication in the **Port number** box.

   d Type the path to the `var` folder for the cluster client in the **Cluster client's var directory** box, or click **Browse** to select a location.

   The `var` folder stores the Exchange DAG client configuration and log files. The VmwareVDPBackupUser account and all nodes in the cluster must have write access to this location.

**NOTE** Select a volume that each server in the DAG can access.

   e Click **Next**.

11 On the **Summary** page, review the settings that you specified in the wizard and click **Configure**.

   The **Progress** page provides the status of the configuration.

12 On the **Results** page, review the information and click **Close**.

## Using the VMware Exchange Backup User Configuration Tool

If the **Launch Exchange Backup User Configuration Utility** checkbox is selected during VMware VDP for Exchange Server Client installation, the VMware Exchange Backup User Configuration Tool starts automatically after installation completes.

### VMwareVDPBackupUser Account

The VDP Microsoft Exchange Server Client requires direct access to the Exchange Server. A special user account, called VMwareVDPBackupUser, is required to provide VDP with appropriate domain and administrator-level permissions. This user account is set up through the VDP Exchange Backup User Configuration Tool, which runs by default after the VDP Microsoft Exchange Server Client installation.

VMwareVDPBackupUser is configured with the following:

- The user account is added and activated for the appropriate Active Directory, Exchange, and group accounts. The user account is added to the following groups:
  - Backup Operators
  - Domain Users
  - Domain Admin (for Exchange Server 2007)
  - Exchange Servers
  - Exchange Organization Management (in Microsoft Exchange Security Groups) for Exchange Servers 2010 and 2013
  - Exchange Organization Administrators for Exchange Server 2007
- A mailbox is created, activated, and tested for the user account.
- A user account is set up and activated in the Exchange domain, and then on each Exchange Server running the VDP Microsoft Exchange Server Client. VDP Backup services must be configured to use the VMwareVDPBackupUser account.

### Prerequisites

- Before using VDP, you must install and configure the VDP appliance described in "VDP Installation and Configuration" on page 21.
- The **Launch Exchange Backup User Configuration Utility** checkbox must be selected during VMware VDP for Exchange Server Client installation. If the user does not select the checkbox at installation, the user can start the VMware Exchange Backup User Configuration Tool manually from the following location:

  *x*:\program files\avp\bin\vmbackupusercfg.exe
- .NET 4.0 must be installed on the Exchange Server:

### Procedure

1. In the VMware Exchange Backup User Configuration Tool, select **New user** or **Existing user** depending on the status of the user you are configuring.
2. In the **User Name** field, type a username for the VMwareVDPBackupUser Account. You can use the default name of VMwareVDPBackupUser if needed.
3. In the **Password** field, type a password for the account.
4. In the **Confirm password** field, re-type the password.
5. In the **Exchange Server** field, select the name of the Exchange Server that the VDP Microsoft Exchange Server Client was installed on.
6. In the **Storage group** field (only active for Exchange Server 2007), select the Storage group name.
7. In the **Mailbox store** field, select the mailbox database for the VMwareVDPBackupUser account.
8. Click **Check** to test the new user settings. If the user does not exist in Active Directory, the check fails.
9. Click **Configure Services**.
10. The message log lists a set of tests that have passed or were successful. If all of the check test pass, click **Close**.

## Manually Configuring the VDP Backup Service

If you have already run the VDP Exchange Backup User Configuration Tool, the VMwareVDPBackupUser account is created. The following steps are used to manually configure the VMwareVDPBackupUser account to run the VDP Backup Service.

- Before using VDP, you must install and configure the VDP appliance described in "VDP Installation and Configuration" on page 21.
- The VMwareVDPBackupUser account has been created through the Launch Exchange Backup User Configuration Utility.

### Procedure

1. Log in to the Exchange Server as VMwareVDPBackupUser or as another user with Administrative rights.
2. Start the Services application by selecting **Start** > **Programs** > **Administrative Tools** > **Services**.
3. From the Services window, right-click **Backup Agent** in the Services list, and select **Properties**.
4. From the Backup Agent Properties dialog box, click the **Logon** tab.
5. Select the **This account** button and specify the username created by the VDP Exchange Backup User Configuration Tool (VMwareVDPBackupUser by default).
6. Type the password for the VMwareVDPBackupUser account in the **Password** and **Confirm password** fields, and then click **OK**.
7. Right-click the Backup Agent service in the Services list and select **Restart**.

**NOTE** If GLR is installed, you must repeat this procedure on the "VMware VDP Exchange GLR Service" as well.

## Creating Backup Jobs for Microsoft Exchange Servers

Microsoft Exchange Server clients must have the VMware VDP Microsoft Exchange Server Client installed in order to be available for backup. See "Installing VDP for Exchange Server Client" on page 174 for additional information on client installation.

1. On the vSphere web client, select the **Backup** tab.
2. In the **Backup** tab, click **Backup Job Actions**, and select **New** to start the **Create a new backup job** wizard.
3. On the Job Type page of the wizard, select **Applications**. This option lets you back up the full server or selected databases.

   The job type that you choose determines the options that you can select from here on. Based on your selection, follow the instructions in the appropriate section below.

### Backing up Applications

If you select **Applications** on the Job Type page, you can choose to back up application servers or individual databases.

1. On the **Job Type** page of the Create a new backup job wizard, select **Applications** and click **Next**.
2. On the **Data Type** page, select one of the following options and click **Next**:
   - **Full Server**: This option lets you back up entire applications servers.
   - **Selected Databases**: This option lets you back up individual application server databases.
3. On the **Backup Sources** page, click the arrow beside **Microsoft Exchange Servers(s)** to expand the list.

4 Do one of the following:

**NOTE** Best practice is to select only one Exchange Server per backup job.

- If you chose to back up a full server, select the checkbox next to the Exchange Server that you want to back up, and then click **Next**.
- If you chose to back up selected databases, click the arrow beside an Exchange Server, continue drilling down until you can select the database or storage group that you want to back up, and then click **Next**. Note that if the client is running as a local system account, the user must provide Exchange Administrator credentials to drill down the Exchange server.

  **NOTE** If the backup target is an Exchange 2007 Server, you cannot select an individual database, and must select a storage group instead.

5 On the **Backup Options** page:

a Select a backup type of **Full** or **Incremental**. Incremental backups are automatically promoted to full backups if a full backup does not exist.

**NOTE** If the client is running as a Local System account, the user must provide Exchange Administrator credentials. If not running as a Local System account, credentials are not required.

If you select **Incremental**, you can specify circular logging options. Circular logging allows you to reduce the number of transaction logs resident on the system. For mixed environments where some, but not all, storage groups or databases have circular logging enabled, you can select one of these settings to specify how VDP handles incremental backups.

- **Promote** (default setting): This option promotes an incremental backup to a full backup if any database in the saveset has circular logging enabled. All databases will be backed up whether they have circular logging enabled or not. If one or more databases has circular logging enabled, all databases in the saveset will have any incremental backup promoted to a full backup.
- **Circular**: This option promotes all incremental backups of all databases with circular logging enabled to a full backup and skips any databases that do not have circular logging enabled.
- **Skip**: This option performs an incremental backup of all databases that have circular logging disabled and skips any databases that have circular logging enabled.

b (Available for DAG clusters only) In the **Preferred server order** list box, specify the priority of servers to use to back up the Exchange databases. Specify the server name, not the FQDN. Separate multiple entries with commas. If you do not specify a list, the Exchange VSS plug-in adds all servers in the DAG to the list in alphabetical order.

c (Available for DAG clusters only) From the **Set the preference for what types of database to back up** list, select the type of database to back up:

- Select **Prefer passive** to back up a passive copy of each database, if a healthy passive copy is available. If a healthy passive copy is not available, the VDP appliance backs up the active copy.
- Select **Active** only to back up only the active copy of each database.
- Select **Passive** only to back up only the passive copy of each database. If a healthy passive copy is not available, the backup does not include the database.

6 Select the **Enable multi-stream backup** option if you want to allow the parallel processing of backup jobs using multiple processors. Use the slider bar to select the number of streams to use.

You can use as many as ten streams. Each stream requires a separate processor core. By taking advantage of multi-processors, you can improve backup performance.

7 Click **Next**.

8 On the **Schedule** page, select the backup schedule and start time for the backup job, and then click **Next**.

"Specifying the Backup Schedule" on page 115 provides additional information about how to configure the schedule.

9 On the **Retention Policy** page, select an option for how long to keep the backup, and then click **Next**.

"Setting the Retention Policy" on page 115 provides additional information about how to configure the retention policy.

10 On the **Name** page, type a name for the backup job, and then click **Next**.

11 On the **Ready to Complete** page, review the summary information about the backup job, and then click **Finish**.

12 Click **OK** when you see the confirmation that the backup job was created successfully.

## Restoring Backups of Microsoft Exchange Servers

After you run backups on Microsoft Exchange Servers, you can restore these backups to their original location or to an alternate location.

**CAUTION** The target Microsoft Exchange Server must have the same Microsoft Exchange Server version and service pack as the Exchange Server on which the backup occurred. Otherwise, the restore will fail.

### Procedure

1 On the vSphere web client, select the **Restore** tab.

2 Click the client whose backup you want to restore.

3 Click the backup that you want to restore.

4 To restore the entire contents of the backup, select the **Exchange Information Store** checkbox.

5 After you select all the targets, click the **Restore** button.

**NOTE** If the client is running as a local system account, the user must provide Exchange Administrator credentials. If not running as a local system account, credentials are not required.

6 On the **Select Restore Options** page, the **Restore to Original Location** option is the default setting and cannot be modified.

7 (Optional) For GLR, on the Select Restore Options page, do the following.

- **Restore to Original Location**: Clear this box if you would like to select a different client to create the RDB on or to restore to a different mailbox.
- **Destination Client**: This is the Exchange Server where the RDB will be created and mounted from the backup. This client must be an Exchange Server with the VDP Exchange GLR plug-in installed.
- **Destination Mailbox**: Set to the email address of the mailbox where the selected mailboxes should be restored to.

8 If you want to specify advanced options, click the arrow beside **Advanced options** to expand the list.

Specify the following options:

- **Allow database overwrite**: Forces any existing databases to be overwritten that have the same name or names included in the restore job. When this option is selected, it modifies the "Allow File Restore" flag, which is internal to the Exchange Server.
- **Restore into RSG/RDB**: Restore Storage Groups (RSG) are used in Exchange Server 2007 and Recovery Databases (RDB) are used in Exchange Server 2010 and Exchange Server 2013. RSG/RDB is used to restore to an RSG/RDB instead of to a production database. When you select **Restore into RSG/RDB**, you can configure the following options:
  - **Overwrite existing RSG/RDB**: Overwrites any existing RSG/RDB. Use this option with caution.
  - **RSG/RDB name**: The name of the RSG/RDB that is used for the restore. If a RSG/RDB with the specified name does not already exist, it will be created. If an RSG/RDB with the specified name already exists, use the **Overwrite existing RSG/RDB** option to overwrite it.

- **RSG/RDB database path**: The path where the RSG/RDB database file will be restored (for example, `C:\myrdb`). This is an optional field. The default location is used if this field is left blank.
- **RSG/RDB log path**: The path where the RSG/RDB log file will be restored (for example, `C:\myrdb`). This is an optional field. The default location is used if this field is left blank.

- Select whether to restore but not replay transaction logs by selecting or clearing the **Do not replay the transaction logs** checkbox. When you select this option, you can manually copy additional transaction logs before you mount the database.
- If log file conflicts occur during the restore, use the **Move logs path** box to specify a location to move the existing log files to before the restore. If you do not specify a path for the log files and there is a gap in the transaction log, the restore process automatically moves the current transaction logs to a subfolder named `logs_time_date`. The time and date values are the time and date of the restore. The subfolder is in the transaction log folder for the database or storage group. You can use these logs to analyze the restore operation, if necessary, or apply those logs up to where the failure occurred.
- For Restore Options in Exchange DAG Cluster configurations, refer to "Suspending Replication in a DAG or Cluster" on page 181.

9. On the **Ready to complete** page, review the restore requests and click **Finish**.
10. In the message box that states that the restore has been successfully initiated, click **OK**.
11. Monitor the restore's progress in the **Recent Tasks** panel.

## Suspending Replication in a DAG or Cluster

If you use VDP to restore to an Exchange DAG, VDP requires the replication from active databases or storage groups to passive databases or storage groups is suspended.

Use one of the following methods to suspend the replication from active databases or storage groups to passive databases or storage groups on the DAG:

- When you perform a restore, select the **Automate replication suspension** option. The VDP appliance automatically suspends the replication.
- Before you perform a restore, manually suspend the replication by using the Exchange Management Shell:

  In the Exchange Management Shell, run the following command:

  **suspend-MailboxDatabaseCopy -Identity** *database*\I

  where, *`database`* is the name of the database and *`server`* is the name of the DAG server that contains the passive copy.

## Monitoring Client Activity

You can monitor tasks and events for all the client activity by collecting and analyzing the client logs. The client logs are Microsoft application (MSApp)-related logs. The aggregated client log includes any replication, backup, restore, or automatic backup verification (ABV) job that passed with exceptions or failed. For more information, refer to "Collecting VDP Logs or Diagnostic Information" on page 40.

## Uninstalling the Exchange Server Plug-in

To uninstall the Exchange Server plug-in, use **Programs and Features**.

The Exchange Server GLR plug-in uninstalls automatically when you uninstall the Exchange Server plug-in. You must restart the computer after you uninstall the Exchange Server plug-in.

## Granular Level Recovery on Microsoft Exchange Servers

The VDP Plug-in for Exchange Granular Level Recovery (GLR) mounts a temporary virtual drive on the target server and restores an Exchange Server database or storage group from a backup to a Recovery Database (RDB) or Recovery Storage Group (RSG) on the virtual drive.

Note the following:

- GLR in VDP supports recovery at the mailbox level only. GLR is not supported at the individual item level.
- The backup must be a full backup with the VMware VDP Exchange Server plug-in.
- You can perform GLR operations on backups that contain public folder databases, but you cannot use GLR to browse or restore from the public folder database itself.
- VDP does not support GLR of the Exchange Server mailboxes that contain special characters in the mailbox names.

### GLR System Requirements

The GLR process with the VDP Plug-in for Exchange Server GLR places additional demands on computer hardware and resources beyond the base requirements for backup and restore. Table 17-27 describes the system requirements for the VDP Plug-in for Exchange Server GLR.

**Table 17-27.** GLR System Requirements

| Requirement | Minimum |
|---|---|
| Memory (RAM) | The VDP Plug-in for Exchange Server GLR requires additional memory (RAM). Start with 1 GB of RAM. The required amount of additional or total memory depends on current system performance with existing memory:<br>■ If resources are already strained and performance is slow with regular operations of the VDP Plug-in for Exchange Server VSS, add significantly more memory to support VDP Plug-in for Exchange Server GLR operations.<br>■ If performance is currently adequate with regular operations of the VDP Plug-in for Exchange Server VSS, additional memory may not be required to support VDP Plug-in for Exchange Server GLR operations. |
| Hard drive space | Additional disk space is required to provide a disk staging area to hold writes to the Exchange Server database and log files. |

**NOTE** You also must download the latest MAPI client libraries and CDO 1.2.1 from the Microsoft website, and install the libraries and CDO on each Exchange Server with the Exchange Server GLR plug-in. Some features might not work unless you have the latest versions.

### Multi-Streaming Requirements

Multi-streaming places additional demands on computer hardware and resources beyond the base requirements for the VDP Plug-in for Exchange Server VSS. In addition, there are several configuration recommendations for multi-streaming.

Table 17-28 lists hardware and software recommendations for multi-streaming.

**Table 17-28.** Multi-Streaming Requirements

| Hardware and Software | Recommendations |
|---|---|
| CPU | At least one processor core per stream |
| Memory (RAM) | 48 GB or more |
| Hard disk | 1 disk drive for operating system/Exchange Server installation<br>1 to 2 disk drives or RAID drive group for each Exchange Server database<br>7200 RPM or faster disk drives |
| Network adapter | 1 GB |
| Operating system | Windows Server 2008 SP2 or later |

### Multi-Streaming Exchange Server Configuration Requirements

VDP consumes significantly more CPU during backups with multi-streaming. This additional CPU consumption on an active Exchange Server can impact performance and affect end users.

Ensure that the Exchange Server environment meets the following requirements for multi-streaming:

- Locate each database on a separate physical disk. If possible, locate the database file on one disk and the transaction logs on a separate disk for each database.
- Best results occur when each database or storage group is approximately the same size.

When you specify multi-streaming options for a backup, specify a maximum of one backup stream for each disk in the backup set. For example:

- If you are backing up two databases, with each database on its own disk, specify a maximum of two streams.
- If you are backing up two databases, with each database and its logs on two disks (for a total of four disks), specify a maximum of four streams.

### VSS Requirements

The VDP Plug-in for Exchange Server VSS uses Microsoft Volume Shadow Copy Service (VSS) technology to perform backups. VSS is a framework that enables performance of volume backups while applications on a system continue to write to the volumes.

#### Supported VSS Providers and Writers

The Exchange Server VSS plug-in uses the Microsoft Software Shadow Copy Provider and the following VSS writers:

- Microsoft Exchange Server Store VSS Writer
- Microsoft Exchange Server Replication VSS Writer

The Exchange Server VSS plug-in does not support hardware providers.

#### VSS Snapshot Volume Requirements

The Microsoft VSS framework supports as many as 64 volumes in a VSS snapshot. When you create a dataset or perform an on-demand backup, do not include more than 64 volumes. If you include more than 64 volumes in a snapshot, backup fails and the event log lists the following error:

`VSS_E_MAXIMUM_NUMBER_OF_VOLUMES_REACHED.`

The VSS framework also limits the number of shadow copies to 64 per volume. If the number of shadow copies in a volume exceeds 64, backup fails and the event log lists the following error:

`VSS_E_MAXIMUM_NUMBER_OF_SNAPSHOTS_REACHED.`

### Enabling GLR Log Files before you Perform GLR

GLR log files enable you to trace and debug granular level restore. Perform the following steps to enable the GLR log files:

1. Use a text editor to create the command file in the `C:\Program Files\avp\var` folder, where `C:\Program Files\avp` is the installation folder.
2. Save and close the command file.

Table 17-29 lists the VDP log files that you can enable to trace and debug GLR.

**Table 17-29.** VDP Log Files

| Log File | Contents | Flag to Enable Debugging | Command File for Flags |
|---|---|---|---|
| `Axionfs.log` | Trace and debugging information for AvFS file system calls. | `--debug` `x19=327680` | `axionfs.cmd` |
| `avmapi.log` | Trace and debugging information for MAPI calls. | `--debug` | `avmapi.cmd` |
| `avexglr_plugin.log` | Trace and debugging information for RDB and RSG creation, mount, browse, and restore. | `--debug` | `avexchglr.cmd` |
| `Aveexchglrsvc.log` | Trace and debugging information for RDB and RSG creation, mount, browse, and restore. | `--debug` | `aveexchglrsvc.cmd` |
| `ps_exec.log` | Trace and debugging information for execution of PowerShell commands on the Exchange Server client. | `--debug` | `ps_exec.cmd` |

### Backing Up Exchange Server Databases

For instructions on how to back up Exchange Server databases, refer to "Backing up Applications" on page 168.

### Restoring Individual Mailboxes

**NOTE** You must manually delete all RDBs on the target Exchange Server (2010 or 2013) before performing granular level recovery. The GLR process automatically deletes RSGs on target Exchange Server 2007 servers.

After you have backed up individual mailboxes, you can browse and extract individual mailboxes from the database to restore. The selected items restore from the VDP appliance to a Recovered Items folder in the original mailbox. You can then browse and select the items to keep. For instructions on how to restore a backup at the granular level, refer to "Restoring Backups of Microsoft Exchange Servers" on page 180.

The virtual drive automatically dismounts and the RDB or RSG is deleted from the target server for GLR after the restore completes. Exchange Server 2007 is automatically deleted. For Exchange Servers 2010 and 2013, you must automatically delete before you perform GLR.

### Procedure

1 On the vSphere web client, select the **Restore** tab.

2 Select the backup to restore and click the **Restore** icon to start the Restore backup wizard.

  The Select backup page appears.

3 Click the client whose backup you want to restore.

4 Click the backup that you want to restore.

5 Select to restore at any level in the hierarchy. The restore is considered GLR if the restore targets are at the mailbox level.

  - For Microsoft Exchange 2007 servers, the backup hierarchy is as follows:

    Exchange server name - > Backup date - > Exchange Information Store - > Storage groups - > Databases - > Mailboxes.

  - For Microsoft Exchange 2010 or 2013 servers, the backup hierarchy is as follows:

    Exchange server name - > Backup date - > Exchange Information Store - > Databases - > Mailboxes.

6 Once all the targets have been selected, click the **Restore** button.

**NOTE** If the client is running as a local system account, the user must provide Exchange Administrator credentials. If the client is not running as a local system account, credentials are not required.

7 On the Ready to complete page, review the restore requests, and then click **Finish**.

8 Click **OK** when you see the message telling you that your restore was successfully initiated.

9 Monitor the restore's progress in the Recent Tasks pane.

# Backing Up and Restoring Microsoft SharePoint Servers

VDP supports the backup and restore of Microsoft SharePoint Servers. Only one SharePoint Server system per farm is currently supported.

**NOTE** Restoring only the SharePoint Server Configuration or Administrator's database can corrupt the SharePoint Server application. Therefore, if you are not restoring only content databases, you must restore the entire backup.

## Hardware Requirements

Table 17-30 lists the hardware requirements for the Microsoft SharePoint Server.

**Table 17-30.** Hardware Requirements for Microsoft SharePoint Server

| Requirement | Minimum |
|---|---|
| Memory (RAM) | 2 GB |
| File systems | NTFS |

## Supported Microsoft SharePoint Servers

VDP supports the following versions of Microsoft SharePoint Server:

- SharePoint Server 2007 SP2 or later:
  - Windows Server 2008 R2
  - Windows Server 2008
- SharePoint Server 2010, 2010 SP1:
  - Windows Server 2008 SP2
  - Windows Server 2008 R2
- SharePoint Server 2013:
  - Windows Server 2012
  - Windows Server 2008 R2 SP1 or later
- SharePoint Server 2013 SP1:
  - Windows Server 2012 R2

## Installing VDP for SharePoint Server Client

Install the VDP client plug-in on each SharePoint Server in the farm. A SharePoint farm is a collection of SharePoint Servers that work together to provide a set of basic SharePoint Server services that support a single site.

The following figure illustrates SharePoint topologies for small, medium, and large farms:

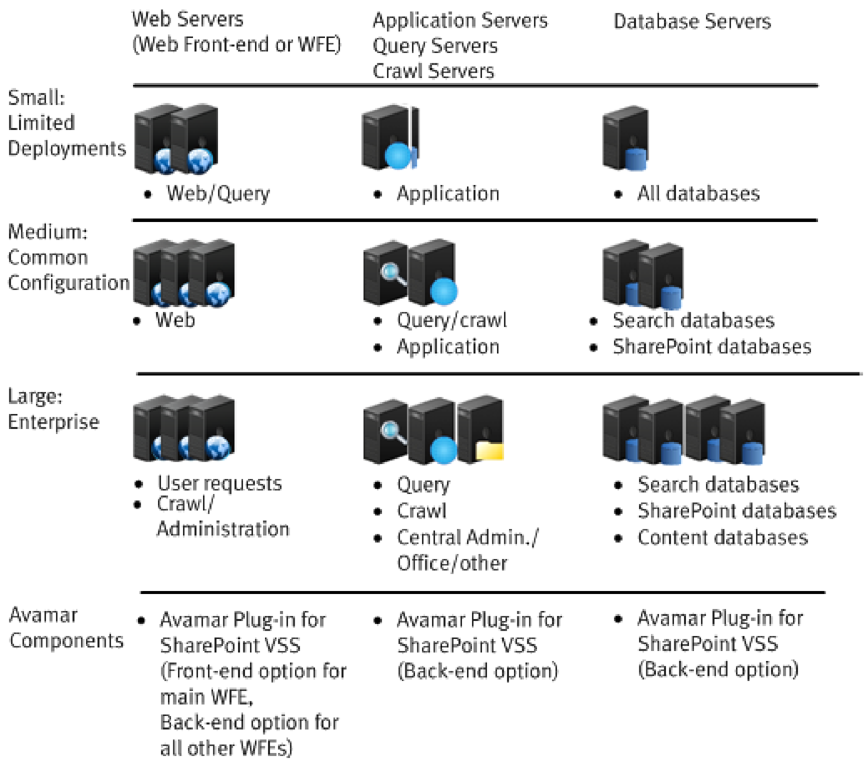



**Figure 17-11.** SharePoint topologies

### Limitations

The SharePoint Server VSS Writer must run with the SharePoint Server Farm Administrator account.

### Prerequisites

- The VDP appliance must be installed and configured as described in "VDP Installation and Configuration" on page 21.
- You must have local administrator privileges to each SharePoint Server.

### Procedure

1. On each SharePoint Server client, access the vSphere web client:

   **https://**<*IP_address_vCenter_Server*>**:9443/vsphere-client/**

2. On the **Credentials** page, enter an administrative vCenter username and password and click **Login**.
3. On the vSphere web client, select **VDP**.
4. On the **Welcome to VDP** page, select the VDP appliance and click **Connect**.
5. Click the **Configuration** tab.
6. Under **Client Downloads**, click **Microsoft SharePoint Server 64 bit**. Depending on your browser, you can save the `.msi` file or run it.

   The VMware VDP for SharePoint Server Setup wizard starts.

7 Click **Next**.

8 On the **End-User License Agreement** page, read the license and if acceptable, click **I accept the terms in the License Agreement** and click **Next**.

9 On the appliance **Registration Information** page, type in the IP address or the name of the VDP appliance that will back up the SharePoint Server. Click **Next**.

10 On the **Ready to install VMware VDP for SharePoint Server** page, click **Install**.

11 During installation, select whether the server on which you are installing will be the primary back up server (front end) or another member server of the SharePoint Server farm (back end).

**NOTE** The front-end server can only be installed on one farm server, and it must be a Web front end or an application server.

12 On the **Completed the VMware VDP for SharePoint Server Setup Wizard** page, click **Finish**.

## Creating Backup Jobs for Microsoft SharePoint Servers

After you have installed the VDP for SharePoint Server client, SharePoint VM clients are available for backup.

1 On the vSphere web client, click the **Backup** tab.

2 In the **Backup** tab, click **Backup Job Actions**, and then select **New** to start the **Create a new backup job** wizard.

3 On the **Job Type** page of the wizard, select **Applications**. This option backs up the SharePoint Server application running on the virtual machine.

4 Click **Next** and follow the instructions below.

### Backing Up Applications

You can only back up the entire SharePoint Server application server. Backing up individual databases is not supported in the current release.

1 On the **Data Type** page, select **Full Server** and click **Next**.

2 On the **Backup Sources** page:

  a Click the arrow beside **Microsoft SharePoint Servers(s)** to expand the list.

  b Select the checkbox next to the SharePoint Server that you want to back up.

  c Click **Next**.

3 On the **Backup Options** page, scroll down, if necessary, to locate the SharePoint Server you have installed as the front end:

  a In the **Farm Administrator Credentials** section, type the administrator **Login ID** and **Login password**.

  b Select the **Enable multi-stream backup** option to allow multiple threads of execution during the backup. Use the slider bar to select the number of streams to use, and then select **Database** or **Volume** in the **Group by** menu.

  c Click **Next**.

4 On the **Schedule** page, select the backup schedule and start time for the backup job, and then click **Next**.

5 On the **Retention Policy** page, select an options for how long to retain the backup, and then click **Next**.

6 On the **Name** page, type a name for the backup job, and then click **Next**.

7 On the **Ready to Complete** page, verify your selections.

If the settings are correct, click **Finish**. If the settings are not correct, click **Back** to go to the appropriate page in the wizard.

## Restoring Backups of Microsoft SharePoint Servers

After you run backups on SharePoint servers, you can restore these backups to their original location or to an alternate location.

1 On the vSphere web client, click the **Restore** tab.

2 Select the backup to restore and click the **Restore** icon.

3 On the **Select backup** page, select the backup job that you want to restore. While you can select multiple SharePoint Servers, you can only select one restore point for each server. Make (or confirm) the backup jobs for restoration and click **Next**.

4 On the **Select Restore Options** page:

   a Perform one of the following steps:

   - Leave the **Restore to Original Location** option selected (the default setting) to restore the backup to its original location.
   - Clear the **Restore to Original Location** option to restore the backup to an alternate location. Click **Choose** to select a destination and type the full Windows path to the location on the destination where the backup will be restored.

   b To specify advanced options, click the arrow beside **Advanced options** to expand the list.

   - **Login ID**: Type the SharePoint Server Administrator's login ID. The format is **DOMAIN\user**.
   - **Login Password**: Type the Farm Administrator's password used to log in to the destination client.
   - **Application Pool** (optional): For SharePoint Server 2013, type the name of an existing application pool to which the Search Service Application will be restored.
   - **Encryption method**: Select an encryption method from the list.
   - **Advanced options** (Support Only): Do not type anything in this box. It is for EMC Support use only.

   c Click **Next**.

5 On the **Ready to complete** page, verify your selections. If they are correct, click **Finish**. If the settings are not correct, click **Back** to create the correct configuration.

6 In the message box that states that the restore has been successfully initiated, click **OK**.

7 Monitor the restore's progress in the **Recent Tasks** panel.

## Monitoring Client Activity

You can monitor tasks and events for all the client activity by collecting and analyzing the client logs. The client logs are Microsoft application (MSApp)-related logs. The aggregated client log includes any replication, backup, restore, or automatic backup verification (ABV) job that passed with exceptions or failed. For more information, refer to "Collecting VDP Logs or Diagnostic Information" on page 40.

## Uninstalling the VDP Plug-in for SharePoint Server

To uninstall the VDP Plug-in for SharePoint Server on a Windows Server 2008 or Windows Server 2012 installation, use **Programs and Features**.

# VDP Disaster Recovery 18

This chapter includes the following topic:


- "Basic Disaster Recovery" on page 190

## Basic Disaster Recovery

In case of a disaster, you can recover the data or VDP appliance by using different methods.

- Using the VDP Checkpoints

  VDP is robust in its ability to store and manage backups. In the event of failure, the first course of action is to roll back to a known validated checkpoint (see"Rolling Back an Appliance" on page 43).

- Using Data Domain Checkpoint Copy

  If you use Data domain integration in the environment, you can use the checkpoint copy feature to restore the failed VDP appliance. The checkpoint copy takes a backup of the VDP checkpoints on Data Domain. You can use these checkpoints to restore the VDP appliance in case of failure. "Restoring Avamar Checkpoint backups from Data Domain systems" on page 92 provides information.

- Detaching and Reattaching VDP Storage

  In case of only appliance OS disk failure or corruption, you can deploy a new VDP and reattach data disks of the failed VDP. "Detaching and Reattaching Storage" on page 75 provides information.

- Using Replication

  You can use the Replication feature of VDP to recover the data in case of a disaster. You can replicate the backups to a remote VDP appliance. This enables you to recover the data from the secondary VDP appliance in case of a disaster on the primary VDP appliance. "Replication" on page 139 provides information.

# VDP Port Usage

A

Table A-31 lists the ports that VDP uses.

**Table A-31.** VDP port usage

| Product | Port | Protocol | Source | Destination | Purpose |
|---|---|---|---|---|---|
| VDP | 7 | TCP | VDP | Data Domain system | ECHO service<br>Required to register a Data Domain system |
| VDP | 21 | TCP | VDP | FTP server | FTP |
| VDP | 22 | TCP | User | VDP | Secure Shell (SSH) access for debugging |
| VDP | 23 | TCP | VDP | Internal | Internal |
| VDP | 25 | TCP | VDP | EMC Customer Support | SFTP |
| VDP | 53 | TCP/UDP | VDP | DNS server | Required for name resolution |
| VDP | 67 | UDP | DHCP server | VDP | DHCP |
| VDP | 68 | UDP | VDP | DHCP server | DHCP |
| VDP | 69 | TCP | Internal switch | VDP | TFTP |
| VDP | 80 | TCP | VDP | vCenter | http for licensing |
| VDP | 88 | TCP/UDP | VDP | Key Distribution Center (KDC) | SMTP |
| VDP | 111 | TCP/UDP | VDP | vSphere Host | Access to RPC port mapper functionality<br>Required when backups are stored on a Data Domain system |
| VDP | 123 | TCP/UDP | NTP time servers | VDP | NTP |
| VDP | 161 | TCP | Data Domain system | VDP | Getter/setter port for SNMP objects from a Data Domain system<br>Required to store Avamar client backups on a Data Domain system |
| VDP | 162 | TCP/UDP | VDP | SNMP Agent | SNMP traps |
| VDP | 163 | TCP | VDP | SNMP service on Data Domain | MC SNMP Manager |
| VDP | 389 | TCP/UDP | VDP | LDAP | LDAP |
| VDP | 443 | TCP | VDP | vCenter or SSO | https |
| VDP | 464 | TCP/UDP | VDP | vCenter | Kerberos |

**Table A-31.** VDP port usage

| Product | Port | Protocol | Source | Destination | Purpose |
|---|---|---|---|---|---|
| VDP | 514 | UDP | Utility node or single-node server | VDP | syslog |
| VDP | 703 | TCP | Avamar server nodes | VDP | AKM service |
| VDP | 902 | TCP | VDP | VMware ESX server proxy | VMware ESX server proxy service |
| VDP | 2049 | TCP/UDP | VDP | NFS daemon on Data Domain system | NFS daemon on Data Domain system |
| VDP | 2052 | TCP/UDP | VDP | NFS mount process on Data Domain system | NFS mount process on Data Domain system |
| VDP | 2888 | TCP | Avamar Extended Retention Media Access Node | VDP | AVDTO<br>Required to support Avamar Extended Retention |
| VDP | 3008 | TCP | VDP | Active archive service on Data Domain system | Active archive service on Data Domain system |
| VDP | 5480 | TCP | VDP | External proxy | External proxy |
| VDP | 5489 | TCP | VDP | Avamar Proxy | Proxy service |
| VDP | 5555 | TCP | ■ Utility node running Avamar Client Manager<br>■ PostgreSQL administrator client computers | VDP | PostgreSQL administrator server |
| VDP | 5568 | TCP | Avamar Extended Retention Media Access Node | VDP | PostgreSQL |
| VDP | 5671 | TCP | ■ Localhost<br>■ Other Avamar utility nodes<br>■ Avamar Extended Retention computers<br>■ EMC Backup and Recovery Manager computers | VDP | RabbitMQ |
| VDP | 6667 | TCP | Avamar Extended Retention Media Access Node | VDP | Archive Service Event |
| VDP | 7000 | TCP | Avamar Extended Retention Media Access Node | VDP | Apche Tomcat |
| VDP | 7443 | TCP | Avamar Extended Retention Media Access Node | VDP | Apche Tomcat |
| VDP | 7780 | TCP | Avamar Administrator Management console | VDP | RMI |

**Table A-31.** VDP port usage

| Product | Port | Protocol | Source | Destination | Purpose |
|---|---|---|---|---|---|
| VDP | 7781 | TCP | Avamar Administrator Management console | VDP | RMI |
| VDP | 7937 - 9936 | TCP/UDP | VDP | NetWorker server | NetWorker service |
| VDP | 8080 | TCP | VDP | NetWorker server | Communication port for NetWorker server |
| VDP | 8105 | TCP | Client computers | VDP | Apache Tomcat |
| VDP | 8109 | TCP | Client computers | VDP | Apache Tomcat |
| VDP | 8181 | TCP | Client computers | VDP | HTTP |
| VDP | 8444 | TCP | Web browser clients | VDP | HTTPS |
| VDP | 8505 | TCP | Utility node or single-node server | VDP | Apache Tomcat |
| VDP | 8509 | TCP | Utility node or single-node server | VDP | Apache Tomcat AJP Connector |
| VDP | 8543 | TCP | vSphere web client | VDP | Redirect for Tomcat |
| VDP | 8580 | TCP | vCenter | VDP | VDP Downloader |
| VDP | 9443 | TCP | vCenter | VDP | VDP Web Services |
| VDP | 19000 | TCP/UDP | Avamar server nodes | VDP | GSAN |
| VDP | 19500 | TCP/UDP | Avamar server nodes | VDP | GSAN |
| VDP | 20000 | TCP/UDP | Avamar server nodes | VDP | GSAN |
| VDP | 20500 | TCP/UDP | Avamar server nodes | VDP | GSAN |
| VDP | 25000 | TCP/UDP | Avamar server nodes | VDP | GSAN |
| VDP | 25500 | TCP/UDP | Avamar server nodes | VDP | GSAN |
| VDP | 26000 | TCP/UDP | Avamar server nodes | VDP | GSAN |
| VDP | 26500 | TCP/UDP | Avamar server nodes | VDP | GSAN |
| VDP | 27000 | TCP | VDP | vCenter | Licensing communication |
| Note: VDP requires that port 27000 be open for licensing purposes. Port 27000 is not required for vCenter communication. | | | | | |
| VDP | 27500 | TCP | Avamar server nodes or Replicator source node | VDP | Avamar server |
| VDP | 28001 | TCP | Client computers, VMWare proxy | VDP | Avamar server CLI, MCS |

**Table A-31.** VDP port usage

| Product | Port | Protocol | Source | Destination | Purpose |
|---|---|---|---|---|---|
| VDP | 28001 | TCP | Replication Source (VDP) | Replication Destination (VDP or Avamar) | Bidirectional communication between avagent and MCS services on the replication source server and the replication destination server to permit authentication key exchange. |
| VDP | 28002 - 28011 | TCP | Avamar Extended Retention Media Access Node | VDP | Support for Avamar Extended Retention |
| VDP | 28009 | TCP | VMware proxy | VDP | avagent |
| VDP | 29000 | TCP | Client computers | VDP | Avamar server SSL |
| VDP | 30001 | TCP | Client computers<br>VMware proxy | VDP | MCS |
| VDP | 30002 | TCP | VDP | Client computers | Required for client or avagent connections |
| VDP | 30003 | TCP | Client computers | VDP | MCS |
| VDP | 30102 - 30109 | TCP | VMware proxy | VDP | avagent |
| VDP | 61617 | TCP | Avamar Extended Retention Media Access Node | VDP | Apache ActiveMQ SSL |
| VDP | 61619 | TCP | VDP | Client that is running EMC Backup and Recovery Manager | Required to permit communication with EMC Backup and Recovery Manager |

# Minimum Required vCenter User Account Permissions

B

See "User Account Configuration" on page 24 to configure the VDP user or SSO admin user by using the vSphere web client. In high-security environments, you can restrict the vCenter user account permissions required to configure and administer the VDP appliance to all of the following categories:

**Alarms**

- Create
- Modify

**Datastore**

- Allocate space
- Browse datastore
- Configure datastore (for VSAN support)
- Low level file operations
- Move datastore
- Remove datastore
- Remove file
- Rename datastore

**Extension**

- Register extension
- Update extensions

**Folder**

- Create folder

**Global**

- Cancel task
- Disable methods
- Enable methods
- Licenses
- Log event
- Manage custom attributes
- Settings

#### Network

- Assign network
- Configure

#### Resource

- Assign virtual machine to resource pool

#### Sessions

- Validate session

#### Tasks

- Create task
- Update task

#### Virtual machine > Configuration

- Add existing disk
- Add new disk
- Add or remove device
- Advanced
- Change cpu count
- Change resource
- Disk change tracking
- Disk lease
- Extend virtual disk
- Host usb device
- Memory
- Modify device setting
- Raw device
- Reload from path
- Remove disk
- Rename
- Reset guest information
- Set annotation
- Settings
- Swapfile placement
- Upgrade virtual machine compatibility

#### Virtual machine > Guest operations

- Guest Operation Modifications
- Guest Operation Program execution
- Guest Operation Queries

**Virtual machine > Interaction**

- Console interaction
- Device connection
- Guest operating system management by VIX API
- Power off
- Power on
- Reset
- VMware tools install

**Virtual machine > Inventory**

- Create new
- Register
- Remove
- Unregister

**Virtual machine > Provisioning**

- Allow disk access
- Allow read-only disk access
- Allow virtual machine download
- Mark as template

**Virtual machine > Snapshot management**

- Create snapshot
- Remove snapshot
- Revert to snapshot

**vApp**

- Export
- Import
- vApp application configuration

# VDP Troubleshooting

C

This chapter includes the following troubleshooting topics:


- "Troubleshooting the VDP Appliance Installation Issues" on page 200
- "Troubleshooting the VDP Management Issues" on page 200
- "Troubleshooting the VDP Licensing Issues" on page 207
- "Troubleshooting the VDP Backup Issues" on page 201
- "Troubleshooting the VDP Restore Issues" on page 203
- "Troubleshooting the VDP Replication Issues" on page 204
- "Troubleshooting the VDP Integrity Check Issues" on page 205
- "Troubleshooting the Automatic Backup Verification Issues" on page 205
- "Troubleshooting the Restore Client (File Level Recovery) Issues" on page 206
- "Troubleshooting the VDP Appliance Issues" on page 208
- "Troubleshooting the VDP Microsoft Exchange Server Issues" on page 208
- "Troubleshooting the VDP Microsoft SQL Server Issues" on page 210
- "Troubleshooting the VDP Microsoft SharePoint Server Issues" on page 210
- "Troubleshooting the Storage Capacity Issues" on page 210
- "Accessing the VDP Knowledge Base Articles" on page 212

# Troubleshooting the VDP Appliance Installation Issues

If you have problems with the vSphere Data Protection (VDP) appliance installation:

- Confirm that all of the software meets the minimum software requirements. See "Software Requirements" on page 22 for more information.
- Confirm that the hardware meets the minimum hardware requirements. See "System Requirements" on page 23 for more information.
- Confirm that DNS is properly configured for the VDP appliance. See "Preinstallation Configuration" on page 23 for more information.

# Troubleshooting the Installer Package Issues

To find upgrade-related logs that you can use to troubleshoot, refer to the `avinstaller.log.0` file in the Log Bundler.

# Troubleshooting the VDP Management Issues

The following troubleshooting topics describe how to identify and resolve some common issues with managing VDP.

## The VDP appliance is not responding, please try your request again

If you were previously able to connect to VDP and this message appears, check the following:

- Confirm that the username or password that is used to validate VDP to the vCenter server has not changed. Only one user account and password are used for VDP validation. This is configured through the VDP Configure utility. See "vCenter Server Registration" on page 42 for additional information.
- Confirm that the network settings for IP and DNS configuration have not changed since the initial VDP installation. See "DNS Configuration" on page 23 for additional information.

## VDP plug-in is not responding

If the VDP plug-in in the vCenter UI displays the **VDP Plugin is not responding** message, it is mostly because either one or more services of the VDP appliance or the entire VDP appliance is down.

Perform the following steps to resolve this issue:

1. Ping the VDP appliance from a remote machine by running following command:

   **ping** <VDP_appliance_IP_address>

   If the output displays 100% message loss, the VDP appliance is either down or not reachable. Start the VDP appliance by either logging in to the vCenter server that manages this VDP appliance or resolve the network issues.

2. Open the VDP configuration utility by using the following URL:

   **https://**<VDP_appliance_IP_address>**:8543/vdp-configure/**

   If this page is not correctly displayed, the tomcat application server has stopped. Log in to the VDP appliance by using ssh client (putty) and restart the tomcat service by running the following command:

   **emwebapp.sh -restart**

# Troubleshooting the VDP Backup Issues

The following troubleshooting topics describe how to identify and resolve some common issues with backups.

### Loading backup job data

This message can appear for a long time (up to five minutes) when a large number of VMs (~100 VMs) are selected for a single backup job. This message can also apply to lock/unlock, refresh, or delete actions for large jobs. This is expected behavior when very large jobs are selected. This message will resolve itself when the action is completed, which can take up to five minutes.

### Unable to add client {client name} to the VDP appliance while creating backup job {backupjob name}

This error can occur if there is a duplicate client name on the vApp container or the vSphere host. In this case only one backup job is added. Resolve any duplicate client names.

### The following items could not be located and were not selected {client name}

This error can occur when the backed up VMs cannot be located during Edit of a backup job. This is a known issue.

### Backup fails if VDP does not have sufficient datastore capacity

Scheduled backups will fail at 92% complete if there is not sufficient datastore capacity. If the VDP datastore is configured with thin provisioning and maximum capacity has not been reached, add additional storage resources.

### Backup fails if VM is enabled with VMware Fault Tolerance

If a VM has fault tolerance enabled, the backup will fail. This is expected behavior. VDP does not support backing up VMs that have Fault Tolerance enabled.

### When VMs are moved in or out of different cluster groups, associated backup sources may be lost

When hosts are moved into clusters with the option to retain the resource pools and vApps, the containers are recreated, not copied. As a result, it is no longer the same container even though the name is the same. Validate or re-create any backup jobs that protect containers after moving hosts in or out of a cluster.

### After an unexpected shutdown, recent backup jobs and backups are lost

Any time an unexpected shutdown occurs, the VDP appliance uses rollback to the last validated checkpoint. This is expected behavior. See "Rolling Back an Appliance" on page 43 for additional information.

### Storage vMotion operations are not allowed during active backup operations

vSphere vMotion is a feature that enables live migration of the running virtual machines from one physical server to another. vMotion operations are not allowed to run on the VDP appliance during active backup operations. This is expected behavior. Wait until all backup operations have completed before you perform a vMotion operation.

If either an automated vMotion or a storage vMotion starts while backups are performed, sometimes some of the backups fail or snapshot deletion fails. In this case, perform the following steps:

1 Wait till the migration completes.

2 Perform backups of the failed VMs by using **Backup > Backup Job Name > Backup Now**.

**NOTE** Do not perform the backups by using **Reports > Task Failures > Rerun task**.

### Backups fail if certain characters are used in the virtual machine, datastore, folder, or datacenter names

When special characters are used in the virtual machine name, datastore, folder, or datacenter names, the `.vmx` file is not included in the backup. The following is a list of the special characters (in the format of character/escape sequence format) that prevent the `.vmx` file from being backed up:

- & %26
- + %2B
- / %2F
- = %3D
- ? %3F
- % %25
- \ %5C
- ~ %7E
- ] %5D

### Rerun option from the Actions icon, runs backups for all clients, not just for the failed client backups

If one client backup fails in a backup job, and then you use the **Rerun Job** option from the Actions icon on the **Reports** tab to rerun the failed backup, the system runs a backup of all of the clients in the backup job.

To run a backup only for the failed client, select **Backup only out of date source** under **Backup Now** on the **Backup** tab.

### Backup now job starting error

If you perform a **Backup now** job to either a VDP appliance storage or a Data Domain storage that is filled above 95%, an error message, which states that the backup is not possible, appears in the vSphere web client.

## Troubleshooting the VDP Backup Performance Issues

The following troubleshooting topics describe how to identify and resolve some common issues with backup performance.

### Check the performance of the disk subsystem where the backup data set is stored

Run the performance analysis tool which exists on the **https://VDP appliance IP/hostname:8543/vdp-configure** page to check whether the datastores on which VDP backups are stored satisfy minimum read/write value that is recommended for VDP. The following table lists the minimum required read/write values for VDP:

**Table C-32.** Minimum read/write values for VDP

| VDP Appliance Size | Disk Size | Minimum Read Value | Minimum Write Value |
|---|---|---|---|
| 0.5 TB | 256 GB | 60 MB/s | 30 MB/s |
| 1.0 TB | 512 GB | 60 MB/s | 30 MB/s |
| 2.0 TB | 1024 GB | 60 MB/s | 30 MB/s |
| 4.0 TB | 1024 GB | 80 MB/s | 40 MB/s |
| 6.0 TB | 1024 GB | 80 MB/s | 40 MB/s |
| 8.0 TB | 1024 GB | 150 MB/s | 120 MB/s |

### Check whether the backups use Hotadd transport mode, but not NBD transport mode

The HotAdd transport mechanism enables faster backups and restores, and exposes less to the network routing, firewall, and SSL certificate issues. If you use the Network Block Device (NBD) transport mechanism instead of the HotAdd transport mechanism, backups will be slow.

### Check for the network bottlenecks if you use the NBD transport mode

If you cannot use the Hotadd transport mechanism, ensure that there are no network issues in the environment and any other issues on ESXi server network configuration that result in slower backups.

## Troubleshooting the VDP Restore Issues

The following troubleshooting topics describe how to identify and resolve some common issues with restores.

### Restore tab shows a “Loading backups” message and is slow to load

It typically takes two seconds per VM backup to load each of the backups on the **Restore** tab. This is expected behavior.

### Restore tab is slow to load or refresh

If there is a large number of VMs, the **Restore** tab can be slow to load or refresh. In tests with 100 VMs, this can take up to four and a half minutes.

### Powering On of VM fails when backed-up virtual machine was connected to DVS and restored to other ESX

When powering on a virtual machine after you restore it, the vCenter returns a PowerOnFailure error similar to the following:

```
DRS cannot find a host to power on or migrate the virtual machine. Network
interface 'Network adapter 1' uses network '77 d3 02 50 5f 76 ca d7-db f9 42 6c
0f 6f 87 1f', which is not accessible.
```

This error message occurs when the environment where you restore the virtual machine does not have the network connection that was present when you backed up the virtual machine. For instance, a user backs up a virtual machine that is connected to a distributed vSwitch. Then the user restores the virtual machine to an ESX host that is part of a Distributed Resource Scheduler (DRS) cluster. The vSwitch does not exist in the DRS cluster. Powering on the virtual machine, therefore, fails.

To work around this issue, edit the virtual machine settings and set the network connection to the network adaptor. Then power on the virtual machine.

### Disk-level restore does not provide an option to specify target datastores

Disk-level restore to a new location does not provide an option to specify the target datastores for each disk of the virtual machine. Currently, VDP restores all the disks of the virtual machine, including the disks that were skipped during backup, into the specified target datastore.

To work around this issue, specify a target datastore that has enough free space to accommodate all the disks of the virtual machine, including the disks that were skipped during backup.

### Deleted disks are skipped when restoring to original location

If the target VM no longer has the same disk footprint as the original VM that was backed up (if the disks have been removed or deleted from the VM), performing a “Restore to original location” operation, after selecting a restore point timestamp in the Restore pane, will silently fail to restore the missing disk of the VM.

The workaround is to restore the disk to its original location after manually adding the missing disk to the VM. Ensure the disk is the same size as it was when the VM was backed up.

If this workaround fails, restore the disk to a new location to create a new VM. When the restore task completes, detach the restored disks from the new VM and attach them to the required VM.

### Name conflict when restoring a disk to an existing virtual machine

Restoring a disk to an existing VM might fail because of a name conflict. The conflict occurs when an existing .vmdk disk is present in the same datastore where the new disk is being restored.

To work around this issue, rename or remove existing disks that are causing the name conflict.

### No time stamp detail displays for Emergency Restore restore points

When navigating through the restore points in the VDP Configure utility's **Emergency Restore** tab, the user may see no timestamp detail for the restore points. This issue occurs when the time zone is defined as UTC.

Perform the following steps to change the UTC time zone:

1. Open the VDP configuration utility.
2. From the **Configuration** tab, select **Change time zone**.
3. Select a time zone other than the UTC time zone, and click **Save**.
4. Click **Apply Changes**.
5. After the web services restart, log in to the VDP configuration utility and verify the timestamp.

## Troubleshooting the VDP Replication Issues

The following troubleshooting topics describe how to identify and resolve some common issues with replication.

### Last successful and last failed replication information not part of email report

The scheduled and ad hoc email reports that are generated after a replication job completes do not contain information about the last successful replication and the last failed replication in the Replication Jobs Summary.

You cannot obtain information about successful and failed replication jobs from VDP.

### Replication job failure errors

If the destination server is in either a Normal state or a Full Access state, the VDP appliance correctly reports the state of the destination server. If the destination server is in the Admin, Read-Only, or Synchronous state, the VDP appliance reports a miscellaneous error when a replication job fails.

If the destination server is filled above 95%, and you cannot start a **Replicate now** job, a warning message regarding the capacity issue appears. Similarly, you cannot recover replicated backups to either a VDP appliance or a Data Domain storage if the storage is filled above 95% of its capacity. A warning message regarding the capacity issue appears.

With inaccurate reporting of execution errors, the user cannot determine the state of the destination server.

### Replication job progress seems stuck in an incomplete status indefinitely

The VDP task panel delivers a "miscellaneous error" for the replication job and fails to initiate the session with servers for replication. One possible cause of this error is the core management service has stopped running for the target server.

Check the management services and restart the core service if it is not running, and then retry the replication job.

### Multiple replication jobs for different VMs created inside a single job run in series and not in parallel

The replication activity for multiple virtual machines should process in parallel. The sequential behavior occurs only when another replication job with the same clients is already running. In this case, the client replication task waits for the already-running replication job to complete.

To work around this issue, you must add the `com.vmware.vdp.option.replicate.maxtreams` property to the `/usr/local/vdr/etc/vdp-options.properties` file. The default value is 1. After you change the default setting when you add or edit the property, the appliance sets the new value as the maximum concurrent number of processes for a replication job.

## Troubleshooting the VDP Integrity Check Issues

After starting an integrity check there can be a delay of a few seconds before the "VDP: Integrity Check" task shows up in the **Running** tasks tab under Recent Tasks. Similarly, when canceling an integrity check, there can be a delay of several seconds before the task is actually canceled.

In some cases (for example, if the integrity check progress is above 90%), the integrity check may actually complete before being canceled. Even though the integrity check may have completed successfully, the Task Console may still show an error indicating the integrity check was canceled.

If you knew that the Integrity Check Status of the appliance (shown on the **Reports** tab) was "Out of Date" before you started the integrity check, then you can look at the status immediately after canceling the job to see if the cancel operation succeeded. If the Integrity Check Status is "Normal," the check was successful. If the status is "Out of Date," the check was canceled.

## Troubleshooting the Automatic Backup Verification Issues

The following troubleshooting topics describe how to identify and resolve some common issues with Automatic Backup Verification.

### Automatic Backup Verification job fails after renaming the datastore

This error can occur if you rename or move the destination datastore outside of VDP.

Edit the job and select the renamed or moved destination datastore as the new destination. For instructions, refer to "Editing a Backup Verification Job" on page 129.

### One or more ABV jobs fail with a "Failed to create VM" error and leave behind an orphaned VM in the vCenter inventory

ABV has triggered the validation of a restored VM on a host that has an earlier version that is incompatible with the version that was used to create the VM. The fact that this leaves behind an orphaned VM is by design because this is necessary for administrators to properly troubleshoot an issue with restores in this type of situation.

Manually delete or unregister the temporary virtual machines that remain in the vCenter or the datastore inventory.

### When an ABV job is canceled, a VDP_VERIFICATION_XXXX VM is left behind on the destination host

You can work around this issue by looking for and manually removing VDP_VERIFICATION_* virtual machines left over on the destination host after an ABV job cancellation.

### Insufficient space available on datastores for the requested actions

ABV jobs or scheduled jobs fail, sometimes with actual data loss, with error messages in the logs to the effect that datastore actions could not be completed because of insufficient space.

Free up space on datastores and retry the ABV job that failed. Perform ongoing routine space management on datastores to prevent running out of space, especially before scheduled jobs affecting datastore space usage.

# Troubleshooting the Restore Client (File Level Recovery) Issues

The following troubleshooting topics describe how to identify and resolve some common issues with the restore client.

### Login failed, cannot locate vm at *10.100.1.10* in vCenter

This error can occur if you are trying to connect to the Restore Client from a host that has not been backed up by VDP.

Log in to a virtual machine that has been backed up by VDP, and then connect to the Restore Client.

### Login failed, cannot locate vm in vCenter

When you log in to the source virtual machine after you perform a file level restore, the login fails with the following error:

```
Login failed. Cannot locate vm in vCenter
```

This error can occur when you restore a VM image to a new virtual machine without a NIC attached. In this case, the FLR is unable to log in to the source virtual machine for a short period of time after the restore completes. To work around this issue, wait a few minutes after the restore completes before you log in to the source virtual machine.

### Restore operation fails with error code *10007*

If a restore operation fails with error code 10007, “Activity Failed - client error(s)” it may be because you selected a read-only destination (for example, a CD drive) or a removable media device that has no media loaded (for example, a diskette drive).

Try the restore again using a new destination or ensure your destination device is writable.

### During a file level recovery mount, only the last partition is displayed if the VMDK file contains multiple partitions

The restore client does not support extended volumes. This is expected behavior. Perform an image level recovery and manually copy the files needed.

### During an file level recovery mount, unsupported partitions fail to mount

The following disk formats are not supported by the restore client, and it is expected behavior that the restore client mount will fail.

- Unformatted disk
- Extended partitions (Types: 05h, 0Fh, 85h, C5h, D5h)
- Dynamic disks (Windows) / multi-drive partitions (any partition that consists of 2 or more virtual disks)
- Deduplicated NTFS
- Resilient File System (ReFS)
- EFI bootloader
- Encrypted partitions
- Compressed partitions

Perform an image level restore and manually copy the files needed.

### Symbolic links are not displayed in the restore client

The restore client does not support browsing symbolic links.

### After importing VMs, FLR login fails for VMs backed up before the import

To perform an FLR from the restore points that have been imported from the previously-used VDP disks, take at least one backup of the VM by using the new VDP.

### Nested container limitations

When restoring a VMware container that contains other containers (that is, a nested container structure; for example vApp-1 contains several virtual machines, and nested inside vApp-1 is a container called vApp-2, which also contains several virtual machines), the VDP appliance only restores the top-level of the hierarchy. Two interim solutions exist for this limitation:

- Flatten the container structure.
- Add both v-Apps (vApp-1 and vApp-2) as separate container entities so that they can be backed up separately. When restoring, restore vApp-1 first, and then restore vApp-2 into vApp-1.

## Troubleshooting the VDP Licensing Issues

The following troubleshooting topics describe how to identify and resolve some common issues with licensing.

### License violation events are generated even though Host is licensed

When licensing hosts in a cluster, all hosts in the cluster must be licensed. Otherwise, all hosts will be considered in violation of the license agreement. Due to the nature of clusters, there is no way to tell which host a VM belongs to at any given time. Therefore the entire cluster must be covered by valid licenses. It is the administrator's responsibility to assign license keys appropriately.

### Evaluation license expires and system is degraded (unusable)

By design, the VDP appliance shuts down integral services whenever license requirements are not met. This can happen when an evaluation license expires.

Do not use an evaluation appliance for any mission critical data. It is intended only as an opportunity to experiment with the VDP Advanced appliance, and the appliance should be considered throw-away.

If, however, you have since purchased a permanent (non-evaluation) vSphere Data Protection license key, there is a mechanism for recovering the appliance. You must perform the following steps between the one hour intervals at which the system forcibly disables key services.

1 Manually restart all services by using the VDP Configuration Utility.

2 Access the VDP appliance from the vSphere web client.

   a Navigate to the License Assignment portlet on the **Configuration** tab.

   b Add the new permanent license key to the system.

This process upgrades the status of the appliance to permanent and the services should remain operational.

**NOTE** If you have multiple evaluation appliances that need to be recovered in this manner, you must repeat the process above on each appliance. You can simply remove and then re-add the permanent license key to force the update.

Alternatively, you can install the license key once and wait 24 hours for the license audit to run. The license audit process detects the presence of the permanent license key and upgrades the appliance from its evaluation license key. You must manually restart the services by using the VDP Configuration Utility.

# Troubleshooting the VDP Appliance Issues

The following troubleshooting topics describe how to identify and resolve some common issues with the VDP appliance.

### The VMware VDP for Exchange Server Clients or the VMware VDP for SQL Server Clients are no longer registered with the VDP appliance

This problem can occur if the VDP appliance has been renamed, or if the clients were installed and a new checkpoint was not created, and a rollback occurred. To work around this problem, reinstall all of the VDP SQL Server and Exchange Server Clients.

### If a Backup Job contains more than one SQL Server or Exchange Server and the servers have identical database paths, if you select a database in one server instance, and not the other instance with the same path, the second instance with the same path will also be backed up

To resolve this issue, either include only one Exchange Server or SQL Server per backup job or ensure that all of your database paths are unique.

### One or more clients cannot be restored; the client is inactive and there are no comparable clients to which a restore can be made

This problem can occur if a user attempts to invoke the restore wizard without selecting a restore point, or a restore point exists for an unregistered client. In either case, the restore cannot occur when there is an inactive client and no comparable target client exists in the environment.

### VDP appliance Guest OS (Linux) becomes read only

The Linux guest operating system becomes read-only if all of the following symptoms appear:

1 Fails to contact or check the status of services, as shown in the following message:

```
root@ldummyxxx:/usr/local/#: dpnctl status
/bin/chown: changing ownership of `/usr/local/avamar/var/log': Read-only file system
dpnctl: ERROR: running as user "root" - problem opening log file "/usr/local/avamar/var/log/dpnctl.log" (-rw-rw-r--) -
dpnctl: ERROR: traceback on exit:
dpnctl_util: pen_log_file (/usr/local/avamar/bin/dpnctl line YYY)
```

2 Unable to log in to the VDP-Configure UI.

3 Unable to connect to the VDP plug-in from within the vCenter server web client.

The following VMware Knowledgebase article provides more information about file systems that may become read-only: http://kb.vmware.com/kb/51306

Before you contact Technical Support, restart the appliance. This might fix the problem.

# Troubleshooting the VDP Microsoft Exchange Server Issues

The following troubleshooting topics describe how to identify and resolve some common issues with the Exchange Server.

### Unmounted or offline databases are skipped

If a database is unmounted or offline when a backup is performed, the backup skips that database. Generally, this is not an issue because databases that are not mounted are not in production.

### Backups may fail when drive letters and volumes are mixed

If you configure the Exchange Server to point to the same database files through different paths, such as volume `G:\` and `C:\MOUNTPOINT`, then backups may fail.

To avoid this backup failure, configure the Exchange Server databases to point to the database files by using the same path. For example, if you have three databases, DB1, DB2, and DB3, that are at the same location as either drive `G:\` or on `C:\`*`mountpoint`*, then use one, but not both, of the following example paths:

- `G:\DB1`, `G:\DB2`, `G:\DB3`
- `C:\`*`mountpoint`*`\DB1`, `C:\`*`mountpoint`*`\DB2`, `C:\`*`mountpoint`*`\DB3`

### Restore to RDB fails or results in an unusable RDB

VDP uses wait times that Microsoft recommends for RDB stabilization in Exchange Server 2013 before the restore starts. The restore either fails or results in an unusable RDB if the stabilization exceeds the wait time. You can increase the wait time to allow more time for the RDB to stabilize.

To increase the RDB stabilization wait time:

1 Use a text editor to create an `avexchglr.cmd` file in the `C:\Program Files\avp\var` folder, where `C:\Program Files\avp` is the installation folder.

2 Type the following text in the command file:

  **--rdb_stabilize_wait=***n*

  where *`n`* is the wait time in seconds. The default value is 60 seconds.

3 Save and close the file.

### Restore requirements are not met

When you restore an Exchange Server database, the destination Exchange Server must have the same Exchange Server version and service pack as the Exchange Server where the backup was performed.

If the Exchange Server version on the destination and source servers do not match, then the restore fails.

### Log files are moved if gaps are detected

During a normal restore, if a transaction log gap is detected, any existing log files are moved to a folder named logs_TIME_DATE, where TIME and DATE are the time and date of the restore. The folder is created as a subfolder in the transaction log file path of the Exchange Server 2007 storage group or Exchange Server 2010 database. You can use these logs to analyze the restore operation, if necessary, or apply those logs up to when the failure occurred.

### Exchange Server 2007 databases are mounted after restore

Before starting a restore, VDP dismounts all databases in a storage group, even if the databases are not being selected for restore. When the restore completes, VDP attempts to mount all existing databases in the storage group, even if they were not previously mounted. VDP does not attempt to mount databases that do not exist on disk, even if they exist in Active Directory.

### Selective restore of databases from an older backup may fail

If you attempt to restore selected databases from an older backup when newer backups exist, then the restore may fail. If this occurs, delete the `restore.env` file created in the log folder path, along with all the log files in that path, and rerun the restore. Also, check the event logs through the Event Viewer if mounting one or more databases.

# Troubleshooting the VDP Microsoft SQL Server Issues

The following troubleshooting topics describe how to identify and resolve some common issues with the SQL Server.

### Not all databases are visible on SQL Server 2012

This problem can be fixed by adding the Windows system service account to SQL Server administrator group:

1. In SQL Server Management Studio, expand the Security node and then the login node for the instance.
2. Right-click the NT AUTHORITY\SYSTEM account and select **Properties**.

   The Login Properties dialog box appears.
3. Select the **Server Roles** page from the list and select the checkbox next to the sysadmin user.
4. Click **OK**.

### A database backup will fail if it is currently being restored

Microsoft SQL Server does not support backups if the database is in a restore state.

### A SQL Server restore with the tail-log backup option selected fails

This problem can be caused if the last restore was performed after the last full backup. Perform a full backup on the database before restore.

The restore can also fail if the **Tail-log backup** restore option is enabled and the database does not exist or is offline. In this case disable the **Tail-log backup** option.

# Troubleshooting the VDP Microsoft SharePoint Server Issues

The following troubleshooting topics describe how to identify and resolve some common issues with the SharePoint Server.

### SharePoint redirected restore job status shows success even when some databases fail to restore

When an IP address is used instead of a server name, backing up part of a SharePoint farm fails when running a redirected restore, even though the UI indicates success.

Use the server name, not an IP address, when creating an alias for a redirected restore of all or part of a SharePoint farm.

# Troubleshooting the Storage Capacity Issues

Carefully assess your backup storage needs when you evaluate the amount of data to store on the VDP appliance and the Data Domain system. Include estimates from the data that other servers sent to the VDP appliance and the Data Domain system. For example, by replication jobs from other VDP appliances, by sharing the same Data Domain system with other VDP appliances, and so on.

When you deploy a new VDP appliance, the storage quickly fills for the first few weeks because every client that you back up contains unique data. VDP deduplication is more effective when other similar clients have been backed up, or the same clients have been backed up at least once.

After the initial backup, the VDP appliance backs up less unique data during subsequent backups. When the initial backups—the local backups and the backups from other servers exceed the retention periods, you can measure the ability of the system to store about as much new data as it frees each day. This process is referred to as achieving a steady state capacity utilization. The ideal steady state capacity is 80%.

## Monitoring the VDP Local Storage Capacity

You must proactively monitor the VDP storage capacity by using the following methods:

- Use the vSphere web client.

  "Viewing Backup Appliance Configuration" on page 55 provides information.

  You can view the capacity also by performing the steps that "Viewing Information from the Reports Tab" on page 108 describes.

- Use the VDP Configuration utility.

  "Viewing the Storage Configuration" on page 76 provides information.

- Configure and monitor the email notifications and reports.

  "Configuring the Email Notifications and Reports" on page 57 provides information.

## Monitoring the Data Domain Storage Capacity

You can monitor the storage capacity of the Data Domain system by using the following methods:

- Use the vSphere web client.

  "Monitoring by Using the vSphere Web Client" on page 94 provides information.

- Use the VDP configuration utility.

  "Monitoring by Using the VDP Configuration Utility" on page 95 provides information.

- Configure and monitor the email notifications and reports.

  "Configuring the Email Notifications and Reports" on page 57 provides information.

## Monitoring the Capacity Issues

Capacity issues start when the storage of either the VDP appliance or the Data Domain system, or both is nearly full. You can monitor the capacity issues by using the following methods:

- The alarms of the capacity issues are triggered on vCenter. You can view the alarms on either the vSphere client or the vSphere web client.
- You can enable the VDP alarm email notifications to receive emails when a VDP alarm is triggered. "Configuring the Email Notifications and Reports" on page 57 provides information. A VDP alarm includes the alarms of the capacity issues too.

### Impact of the capacity issues

If the storage capacity of either the VDP appliance or the Data Domain system, or both exceed 80%:

- `The` *`VDP/Data Domain`* `storage is nearly full.` yellow alarm is triggered.
- The warning message appears on the alert bar, and in the **Storage Summary** sections on the vSphere web client and in the VDP configuration utility.

If the storage capacity of either the VDP appliance or the Data Domain system, or both exceed 95%:

- `The` *`VDP/Data Domain`* `storage is nearly full.` red alarm is triggered.
- The error message appears on the alert bar, and in the **Storage Summary** sections on the vSphere web client and in the VDP configuration utility.
- You cannot perform Backup now jobs and Recover replicated backups on the VDP appliance and the Data Domain system.

If the storage capacity of either the VDP appliance or the Data Domain system, or both reach 100%:

- `The` *`VDP/Data Domain`* `storage is full.` red alarm is triggered.

- The error message appears on the alert bar, and in the **Storage Summary** sections on the vSphere web client and in the VDP configuration utility.
- The current backups on the VDP appliance and the Data Domain system fail.
- You cannot perform new backups on the VDP appliance and the Data Domain system.
- Replication of data to the VDP appliance and the Data Domain system fails.
- Operations that change information on the Data Domain system fail. The operations include the deletion of checkpoints, active backups, and expired backups during garbage collection. These operations can fail because they involve directory renames, which are not allowed on a full Data Domain system.
- The VDP appliance is in the Admin state.

### General Steps to Free Storage Space

You must perform the following steps to free storage space in advance to avoid reaching the maximum capacity. If the storage capacity exceeds 80%, perform the following steps:

- Do not add virtual machines as backup clients.
- Remove unnecessary restore points.
- Delete unwanted jobs.
- Reassess retention policies to decrease retention period.
- Add VDP appliances and balance the backup jobs among multiple appliances.
- If you have not added a Data Domain system to the VDP appliance, add it and balance the backup jobs between the VDP appliance and the Data Domain system.
- Expand the local VDP storage disk. "VDP Disk Expansion" on page 99 provides information.
- If the Data Domain system reaches its capacity limit, perform the steps that "Reclaiming Storage on a Full Data Domain System" on page 95 describes to reclaim space.

## Accessing the VDP Knowledge Base Articles

Additional troubleshooting information is available through VDP Knowledge Base Articles, which are located at.

http://www.vmware.com/selfservice/microsites/microsite.do

Select **Products** > **VMware vSphere Data Protection** Category > **Troubleshooting**

# Index


## A

Adobe Flash Player **22**
alarms, viewing **61**
appliance
  rolling back **43**
automatic backup verification
  best practices **126**
  described **126**
  troubleshooting **205**
automatic backup verification. *See also backup verification jobs*

## B

backup and recovery
  using changed block tracking **14**
  using datastore **14**
  using file level recovery **14**, **154**
  using Virtual Machine Disk (VMDK) **14**
  using VMware vStorage APIs for Data Protection (VADP) **14**
backup job
  full image **116**
  individual disks **116**
  migrating from VDP to Avamar **121**
  on individual disks **117**
backup schedule **143**
backup verification job
  cloning **129**
  deleting **130**
  disabling **130**
  editing **129**
  monitoring **130**
  running **129**
backup verification jobs
  creating **127**, **130**
  limitations **126**
backup window, editing **57**
backups
  filtering **157**
  increasing number that run **65**
  mounting **157**
  reducing number of concurrent backups **65**
  verifying **130**
best practices
  automatic backup verification **126**
  backups **65**
  Data Domain **81**
  deploying proxies **64**
  emergency restore **44**
  general **26**
  HotAdd **27**, **64**, **66**
  proxies **66**
  replication **142**, **147**
  storage capacity for initial VDP deployment **27**
  supported disk types **22**, **118**
  VDP Appliance deployment **26**

## C

capacity
  insufficient for backups **201**
  monitoring **28**
  requirements **22**
  steady state **27**
  utilization **76**
  VDP, initial deployment **27**
Certificate Authority (CA) **51**
Changed Block Tracking (CBT) **14**, **16**, **26**
collecting logs **40**
compatibility matrix
  replication recovery **149**
  replication source **140**
concurrent backups **65**
configuring Data Domain storage **80**
CPU load **69**
customer experience improvement program **17**, **54**

## D

Data Domain
  adding a system to VDP configuration **85**
  architecture overview **80**
  backups with VDP **89**
  best practices **81**
  capacity monitoring **94**
  changing the max streams value **84**
  client support with VDP integration **80**
  common problems and solutions **97**
  configuration **80**
  deleting from VDP appliance **87**
  limitations in VDP environment **81**
  monitoring from the VDP Appliance **94**
  port usage requirements for VDP communication **83**

pre-integration requirements **82**
preparing the system for VDP integration **84**
reclaiming storage on a full system **95**
selecting a target for backups **90**
server maintenance activity monitoring **91**
data protection
using changed block tracking (CBT) **14**
using datastore **14**
using File Level Recovery (FLR) **14**, **154**
using Virtual Machine Disk (VMDK) **14**
using VMware vStorage APIs for Data Protection (VADP) **14**
deduplication, benefits of **15**
direct to host restore **44**
disk expansion
growing and adding disks **101**
limitation **102**
requirements **100**
with Essentials Plus **103**
disk usage **69**
disks, types supported by VDP **22**, **118**
Distributed Resource Scheduler (DRS) **68**
distributed vSwitch (dvSwitch) troubleshooting **203**
DNS configuration
setting up **23**
verifying **23**


## E


email reports **55**, **57**
emergency restore
automatic host detection **46**
best practices **44**
limitations **44**
performing **44**
troubleshooting restore points **204**
unsupported features **44**
Essential Plus license **100**
ESXi compatibility with vSphere Flash Read Cache **22**
event console, viewing **62**


## F


File Level Recovery (FLR)
benefits **17**
limitations **67**, **154**
overview **14**, **154**
requirement **64**
support **66**
unsupported VMDK configurations **155**
filtering backups **157**
fixed-length data segment **15**


## G


granular level restore (GLR)
log files **183**
Microsoft Exchange 2007 servers **184**
Microsoft Exchange 2010 or 2013 servers **184**
troubleshooting **183**


## H


hardware versions
migrating **22**
upgrading **22**
HotAdd **22**, **27**, **64**, **66**, **68**, **154**


## I


image-level backups **16**
individual disk backups
procedure **118**
supported disk types **118**
individual disks
creating a backup job **16**, **117**
impact when migrating **119**
installation requirements **22**
installing vSphere Data Protection (VDP)
software requirements **22**
integrity checks, running **59**
internal proxy, disabling **68**


## J


job details tab **110**


## K


knowledge base, accessing articles **212**


## L


license audit **207**
license key
permanent **207**
requirement **15**
licensing
Essential Plus **100**
troubleshooting evaluation license expiration **207**
log bundle, file name of **40**, **200**
log collection **40**
Logical Volume Manager (LVM) **64**, **154**
logs
bundle **40**, **200**
collection **40**
proxy **69**
viewing **59**


## M


man-in-the-middle attacks **51**

master boot record **154**
memory usage **69**
Microsoft Exchange
- backing up applications **178**
- creating backup jobs **178**
- increasing RDB wait time **209**
- installing clients when UAC is enabled **162**
- installing the client **174**
- manually configuring the backup service **178**
- restoring backups **180**
- server options **173**
- using the configuration tool **176**

Microsoft SharePoint
- backing up and restoring servers **185**
- creating backup jobs **187**
- restoring backups **188**

Microsoft SQL
- backing up and restoring **163**
- backing up applications **168**
- creating backup jobs **168**
- installing the client **165**
- restoring backups **171**
- server options **163**
- server support **164**

monitoring
- backup verification jobs **130**
- capacity **28**
- Data Domain capacity **94**
- Data Domain from VDP Appliance **94**
- Data Domain maintenance activity **91**
- VDP activity **60**

mount limitations **157**
mounting backups **157**
multi-tenancy, support for **151**

## N
network address translation (NAT) **154**
network block device (NBD) **22**, **27**, **66**
network settings, configuring **42**
node structure, replication **148**
NTFS **155**, **185**, **206**

## O
OVF template file **30**
OVF template file, deploying **29**

## P
performance analysis
- running **103**
- test results **103**

platform product support **16**
port 29000 **150**
provision type **72**
proxies
- adding **67**
- best practices **66**
- deploying **64**
- external support **66**
- health status **69**
- internal, disabling **68**
- managing **65**

proxy logs **69**

## R
replication
- back to source **148**
- best practices **142**, **147**
- cloning a job **147**
- creating a job **143**
- Data Domain as backup target **142**
- deleting a job **148**
- destinations **149**
- editing a job **147**
- enabling or disabling a job **148**
- limitations **143**
- managing destinations **147**
- node structure **148**
- recovery **150**
- running existing jobs immediately **148**
- scheduling and managing jobs **17**, **140**
- selecting backups to replicate **144**
- viewing job status and details **148**

Replication Target Identity (RTI) **140**, **149**
reports tab
- viewing Data Domain status **86**
- viewing information **108**

restore
- to read-only media **206**
- to removable media **206**

restore points, refreshing **46**
Restores to New
- disabling the vMotion feature before running **136**

restoring backups
- direct to host **44**
- manually **133**, **135**, **136**
- to new location **134**, **135**, **137**
- to original location **134**, **135**, **137**
- to SCSI disk ID **134**, **135**, **137**
- when snapshots are present **133**

retention policy **117**, **143**
rolling back an appliance **43**

## S
security
- Certificate Authority (CA) **51**

login attempts **54**
man-in-the-middle attacks **51**
SSL certificates **51**
services
backup scheduler **38**
core services **38**
file level restore services **38**
file system services **38**
maintenance services **38**
management services **38**
starting and stopping **39**
status of **39**
show completed activities **159**
shutting down VDP Appliance **62**
software requirements **22**
SSL certificates **51**
SSO admin user **25**
starting VDP Appliance **62**
steady state capacity **27**
storage
attaching existing VDP disks **74**
creating **72**
detaching **75**
import existing **72**
summary **56**
viewing summary **76**
storage summary, viewing **102**
supported disk types **22**, **118**
SVMotion
used for live migration of vmdk files **119**

## T

task failure tab **109**
tasks, viewing **60**
technical support resources **12**
time synchronization errors **24**
troubleshooting
after an unexpected shutdown, recent backups are lost **201**
associated backup sources may be lost **201**
automatic backup verification (ABV) **205**
backup fails if VM is enabled with VMware fault tolerance **201**
backup fails if vSphere Data Protection does not have sufficient datastore capacity **201**
backups are slow to load **203**
backups fail if certain characters are used **202**
file level recovery **206**
granular level restores **183**
guest OS (Linux) becomes read-only **208**
items could not be located **201**
license violation events **207**
loading backup job data **201**
Microsoft Exchange RDB wait time **209**
name conflicts when restoring a disk to an existing VM **204**
powering on virtual machine fails when connected to distributed vSwitch **203**
replication job failure errors **204**
replication jobs **204**
restore operation fails **206**
unable to add client **201**
VDP Appliance is not responding **200**
VDP Exchange backups **208**
VDP SQL backups **210**
VDP SQL restores **210**
vSphere Data Protection Exchange restores **209**
vSphere Data Protection integrity check **205**

## U

upgrading
hardware versions **22**

## V

variable-length data segment **15**
vCenter
user account permissions **195**
vCenter Server
changing hostname, password, and so forth **42**
supported versions **29**
switching VDP Appliances **108**
version requirements **22**
vCenter Server Appliance (VCSA) **16**
VDP Configure Utility **38**
VDP disks
attaching **74**
VDP licensing, troubleshooting **207**
VDP, types of disks not supported **22**, **118**
viewing
alarms **61**
event console **62**
storage summary **76**, **102**
tasks **60**
Virtual Machine Disk (VMDK) **14**
Virtual Volumes (VVOLs) **23**
VMFS heap size **100**
vMotion
disabling the feature before running Restores to New **136**
VMware Tools **24**, **127**
VMware vStorage APIs for Data Protection (VADP) **14**
vSphere Data Protection (VDP)
accessing from command line **106**
accessing knowledge base articles **212**
accessing with vSphere Web Client **106**

changing a configuration **42**
collecting logs **40**
installation **30**
monitoring activity **60**
sizing **22**
software requirements **22**
starting and stopping services **39**
storage capacity **27**
supported configurations **23**
user interface **108**
viewing status of services **38**
vSphere Data Protection (VDP) Appliance
best practices when deploying **26**
corruption **75**
description of **14**
rolling back an appliance **43**
shutdown **62**
startup **62**
switching from vCenter Server **108**
vSphere Flash Read Cache and ESXi compatibility **22**
vSphere HA **72**
vSphere host, supported versions **22**
vSphere Web Client **107**
Adobe Flash Player requirement **22**
supported web browser **22**